AUSTRALIAN
communications
AUGUST 1994 The Networking and Telecommunications Management Magazine $5.50

DIAL M FOR MOVIE
Why the world is watching video-on-demand

UPGRADING AARNET
How should we build an Australian R&D network?

DAYTONA DEBUTS
A comprehensive look at Windows NT 3.5

ETHERNET BRIDGES
Putting the squeeze on WAN costs

TOKEN RING NETS
New technologies mean new design freedom

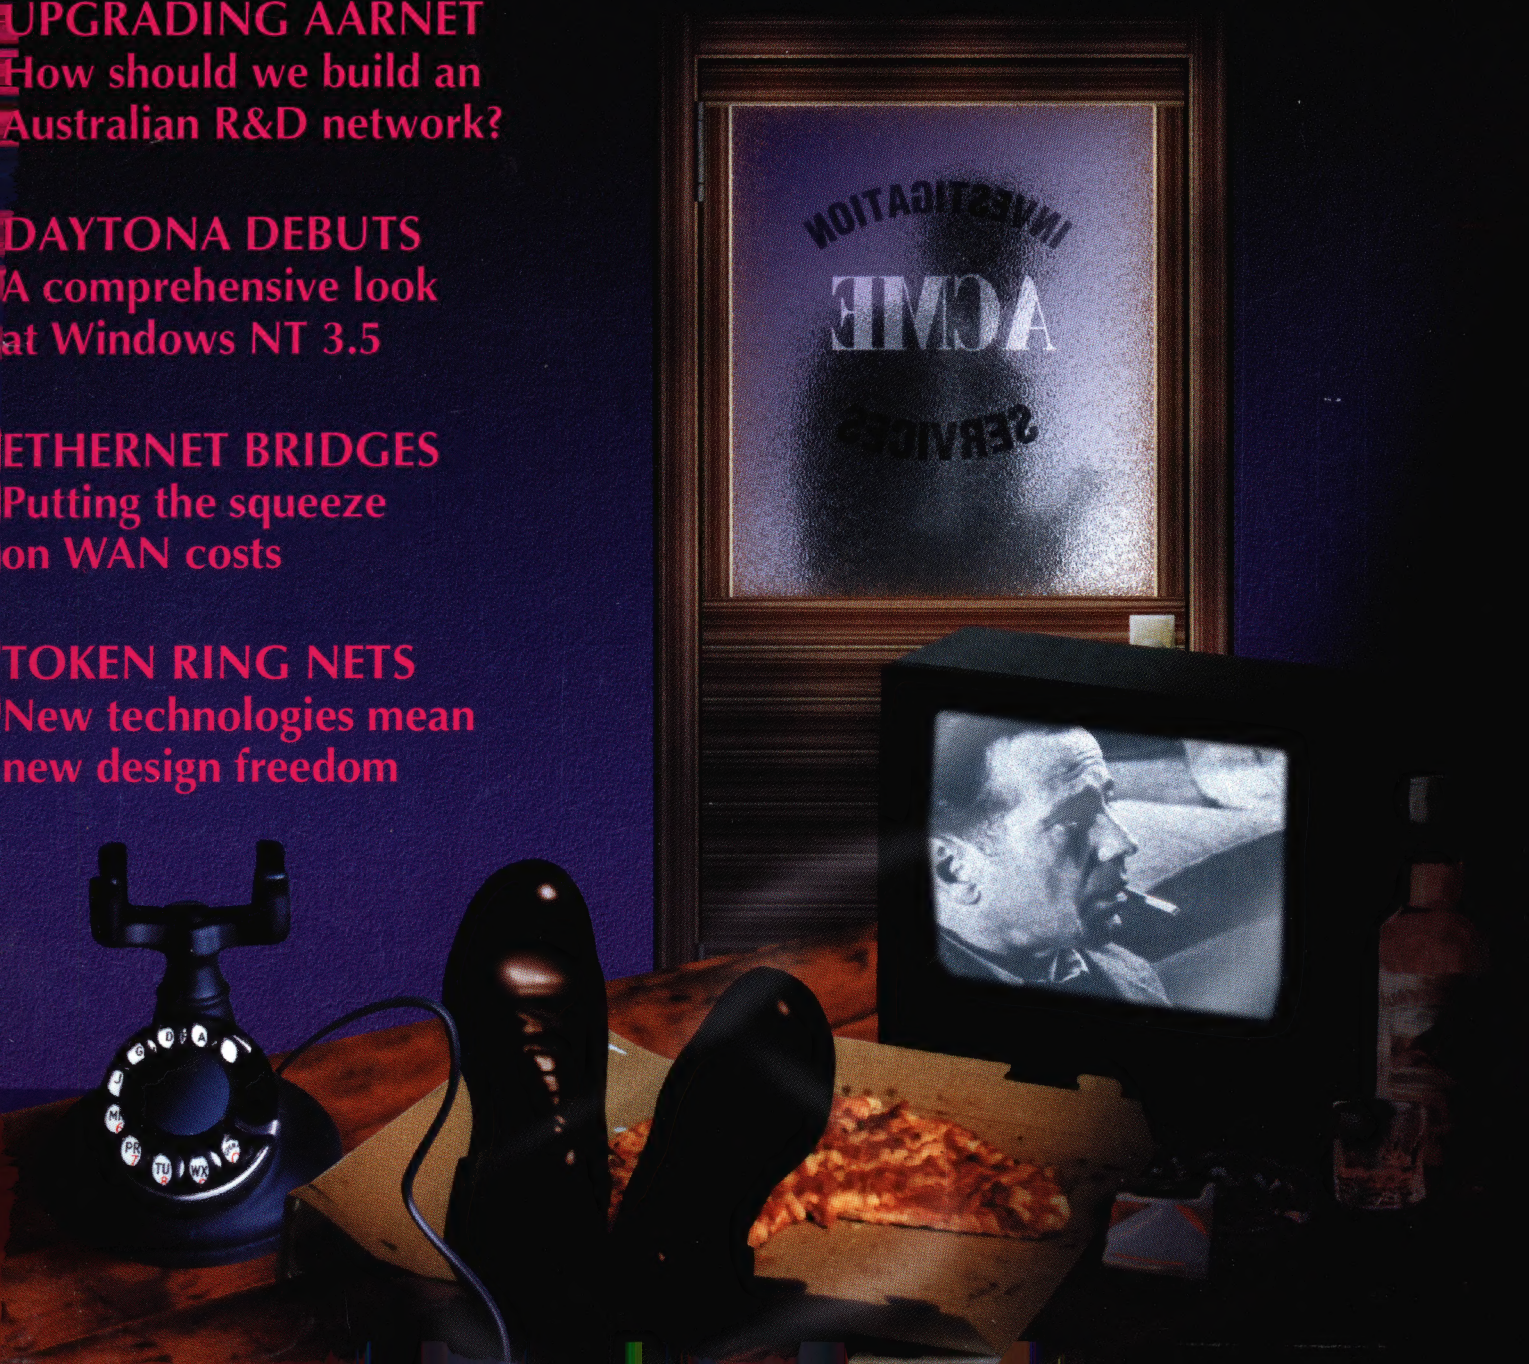

Now that Lotus N
work this well togethe

tes and NetWare

your group can, too.

Every decision your group makes can now be far more informed, much wiser and occur with a speed never before possible.

Because now the world's most powerful workgroup productivity tool fits seamlessly onto the world's most powerful network. Together, they place you at the centre of a whole new class of applications that radically improve strategic business processes - like customer service, account management and product development.

Now you can access, organise and share in minutes information that once took weeks to disseminate. Group members - even if they're widely scattered or only occasionally connected to the network - can make better decisions faster, to make your company more competitive.

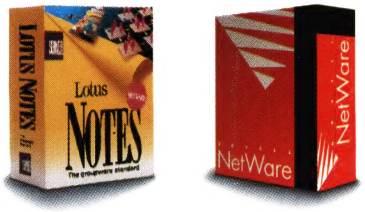

Lotus Notes for NetWare

As a NetWare® Loadable Module™ , Lotus Notes® integrates immediately into your familiar NetWare environment, making it even more powerful.

It's yet another way NetWare acts as a true "plug and play" platform for building advanced applications and services on your information system – right now (without starting from scratch), and well into the 21st century. Leading the way in the networking and open systems market in Australia, Com Tech is the only distributor able to offer this outstanding solution.

Call Com Tech on (02) 317 3088 or 008 263 954 for details of your nearest Lotus Notes partner.

Melbourne (03) 696 0770 Canberra (06) 257 7695

Brisbane (07) 832 2699 Adelaide (08) 271 8600

Perth (09) 322 5605

Lotus
Working Together

NOVELL

PLEASE RUSH ME MORE INFORMATION ON:

☐ Lotus Notes for NetWare ☐ NetWare

Name _____ Title _____

Company _____

Address _____

_____ Postcode _____

Phone _____ Fax _____

My preferred reseller is _____

_____

COM TECH

CT3720

AUSTRALIAN
communications

Cover: Eye For Design

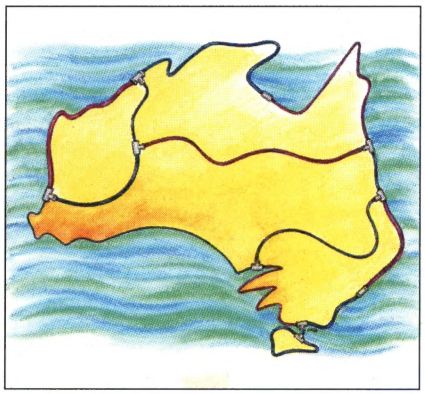

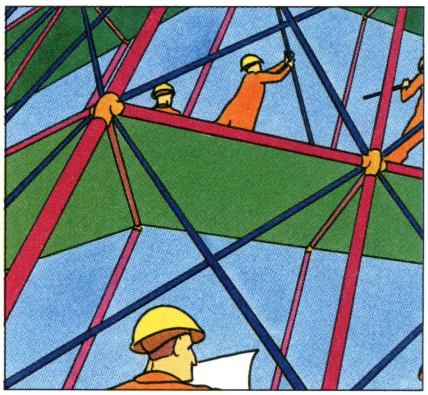

VIDEO-ON DEMAND 65

To go with their dial-a-pizza dinners, couch potatoes around the world will soon have something very special to look forward to: dial-a-movie (a.k.a. video-on-demand). Ultimately promising access to the world's store of movies and other film material, video-on-demand is a new technology now rolling out of the labs and into carefully managed trials. Conceptually nothing more than database access and information retrieval conducted on a one-to-one basis, computer-to-computer, video-on-demand data is delivered as a 90 minute interleaved voice/image digital stream. In practice, things get a little more complicated, as Stewart Fist explains.

AN AUSTRALIAN RDN 75

Efficient data networking is now recognised as a key element of the information economy and as a consequence, super high-speed research networks are now being built in several countries. What should Australia do? The time-honoured thing, of course — appoint a committee. And so it was that in October last year the Federal Government commissioned the Australian Science and Technology Council (ASTEC) to report on Australia's future requirements for Research Data Networks (RDNs). During the preparation of its final report, ASTEC commissioned Melbourne-based Cutler & Company to investigate key RDN issues. We present here the findings of that investigation.

TOKEN RING NETWORKS 85

Token Ring may have lost its networking Holy War with Ethernet but there are still plenty of zealots who refuse to accept that there is a better way. For those die-hards Token Ring will live forever — and with the variety of new generation equipment now coming onto the market, maybe they're right. The arrival of every new and improved technology — first modular Token Ring hubs, now LAN switches, and soon high-speed (100Mbps) hubs and adaptors — increases the temptation to replace the multistation access units that served as the first foundations for Token Ring LANs. To hear vendors tell it, that's exactly what network designers are doing.

AUGUST 1994

ETHERNET BRIDGES 97

Can remote Ethernet vendors live up to their claims about data compression? The answer, happily, is an emphatic 'yes.' The six remote Ethernet bridges put through their paces in the Data Comm Test Lab (the Magnalink Series 3000, Microcom's MBR/6003, Newport Systems Solutions' LAN2LAN/MPR, RND's Remote Compression Bridge, 3Com's NetBuilder and Retix's 4810 LAN Bridge) can compress text by 3:1 or better. And some of this equipment actually deserves the much-overused label 'plug and play.' Considering that leased-line tariffs and management represent the two biggest costs of remote bridging, these findings are very good news indeed for network managers.

ANALYSIS

13 The Privatisation Push
Despite pre-election promises, the privatisation of Telecom is now very much back on the agenda.
14 Stock Watch
15 Telstra Staffing Schemes
16 BellSouth's NZ Blunders
18 There's Money in Resale
18 Upgrade AARNet Now
20 Finally Fixing the FTAs

TECHNOLOGY

23 Microsoft's Daytona Dawn
Daytona is Windows NT 3.5, Microsoft's latest attempt to crash through its final frontier — networking.
24 Market Watch
28 Corporate Connectivity Channels
32 SynOptics Scores at the Soccer
35 Bell Labs' Laser Breakthrough
36 LANexpress Wins Top Marks
37 Technology Update
38 Wellfleet, SynOptics Join Forces
40 Paging Gets PCMCIA
40 Discovering Intelligent Agents
43 Software Views — Frank Henderson
44 LAN Views — Graeme Le Roux
46 WAN Views — Steven Taylor

PRODUCT LEADERS

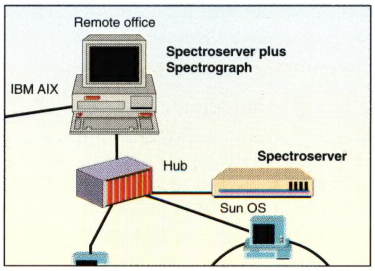

49 Industrial Strength Integration
Fibronics' GigaHUB delivers the goods now and in the future.
50 Making the Office Accessible
3Com's AccessBuilder makes remote network access easy.
51 A Shared Management Burden
Cabletron's Spectrum 3.0 distributes the net management burden.

OPINION

60 Tom Amos explores the notion of services franchising. It works for ice-creams, furniture and hamburgers — but telecommunications?

61 Kevin Morgan, formerly with the CWU and now an independent consultant, discusses Telecom's privatisation. He's in favour.

INTERVIEW

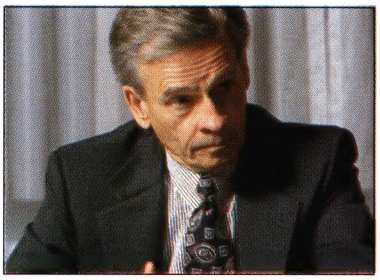

55 Fred A. Vierra
One of four members of the Chief Executive Office for Tele-Communications Inc. (TCI), Fred Vierra is responsible for the company's global portfolio, including its cable TV systems. He discusses communications convergence with Liz Fell.

LEGAL LINE

53 The Post-1997 Policy Review
The Federal Government's recently announced carrier review could result in a new round of sweeping changes to the *Telecommunications Act*. Peter Waters considers the issues which need to be addressed.

DEPARTMENTS

VIEWPOINT 6
UPDATE 8
PRODUCT FILE 107
THE ATUG PAGES 117
CALENDAR 122

Lights, Camera . . . Action!

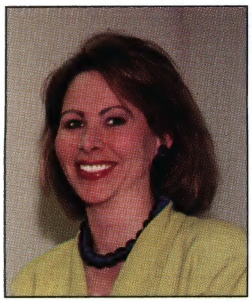

One of the most keenly anticipated products of the much-vaunted information superhighways, video-on-demand systems promise to make it as easy to watch your favourite movie — whether it be as popular as the latest Arnie blockbuster or as obscure as an early Howard Hawks classic — as simple as ordering the pizza to go with it. In practice, however, the technology is still very much on the drawing board, with a number of complex issues still to be nutted out. In our cover story this month, Stewart Fist takes us through the maze of different access and delivery methods, and looks at the kinds of services likely to be offered and — mostly importantly — their cost to users. While there's much speculation on who the big players will be in this highly lucrative market, there's little doubt US giant Tele-Communications will be one of the leaders. In this month's *Interview* on page 55, Liz Fell talks to Fred Vierra, the man responsible for the company's global cable TV interests, about the changing broadcasting and communications marketplace.

All the talk of superhighways has also led some countries to begin to build their own super high-speed research networks, prompting a recent enquiry by the Australian Science and Technology Council (ASTEC) into Australia's future requirements for Research Data Networks. Will AARNet, Australia's link to the global Internet, form the basis of the new Information Age, and in its present form would it be able to handle the projected huge demand for data from the business community? Cutler & Company, the organisation commissioned by ASTEC to report on key RDN issues, presents the results of its findings in 'Building Australia's R&D Network' on page 75.

In the September edition of *Australian Communications* . . .

PCN/PCS

The new wave of personal communications systems promise to revolutionise the way we communicate, but how do these systems work, and how will we use them? Next month Stewart Fist walks us through this new technology.

OVERCOMING CLIENT/SERVER CHAOS

A distributed computing environment generally may improve efficiency, but the client/server model can also have an adverse effect on information flow between departments. We look at ways for network managers to reintegrate the enterprise.

DIAL-UP ROUTERS

Dial-up routers can represent a lifeline for small branch sites that don't have the traffic load to justify their own leased line, but the Data Comm Test Lab reveals many products fail to come up to scratch on ease-of-use.

INTERNATIONAL CASE STUDY: SKF

Sweden's SKF, a global manufacturer with operations in over 40 countries, recently undertook the chore of moving its mission-critical SNA traffic onto a global router backbone. In a special report, we find out how the task was achieved.

AUSTRALIAN communications

EDITOR
Mark Smeaton

CONTRIBUTORS
Tom Amos, Stewart Fist, Liz Fell, Peter Leonard, Peter Waters, Stuart Corner, Graeme Le Roux, Bernard Levy, Bill Bennett, Alan Lloyd, Peter Purton

ASSISTANT EDITOR
Sarah Parkes

PRODUCTION MANAGER
Elisabeth Alexander

ILLUSTRATIONS
Raphaella La Roche, Aldo Venturi
Loui Silvestro

SALES AND MARKETING DIRECTOR
Craig Burkill

DEPUTY ADVERTISING MANAGER
Dawn Switzer

CHIEF FINANCIAL OFFICER
Jennifer Stibbard

Subscriptions
Australian Communications is available by paid subscription for $54 per annum. Subscribers receive eleven issues per annum.

Overseas Rates by Air	1 Year	2 Years
New Zealand, Papua New Guinea	$66	$110
Singapore, Indonesia, Malaysia, Brunei, Pacific Islands	$72	$122
Asia — including Hong Kong, India, Korea, Japan, Taiwan	$78	$133
Europe, North America, Middle East, South America	$92	$168

Publishers
Published and distributed monthly by Ostasun Pty. Ltd. (ACN 003 606 102) under licence from ACP Computer Publications, a division of ACP Publishing Pty. Ltd. (ACN 053 273 546) of 54 Park Street, Sydney NSW 2000. ISSN 0818-9021. **Address:** Level 4, 541 Kent Street, Sydney NSW 2000. **Tel:** (02) 264 2200 **Fax:** (02) 264 2244.

Printed at **Offset Alpine Printing**, Derby & Wetherill Streets, Silverwater NSW 2141.

The solution is to work with a company that knows how to take your equipment and make it work for you. A company that knows what's ahead so you won't get left behind. A company that understands how to build a network with the latest routers, hubs and structured cabling systems that will grow with you. Today, the question isn't just about equipment. But who can you trust with the know-how to bring everything in your network together?

HAVING THE EQUIPMENT IS ONE THING.

KNOWING HOW TO USE IT IS ANOTHER.

Anixter. *Nobody knows networking like we do. It starts by working with the leading manufacturers to design a network perfectly tailored for you. Making sure it stays that way. And supporting it with technical expertise from people who really understand your network. For cabling systems, networking products and the know-how to get you exactly what you need, exactly when you need it, there's only one place to set your sights:*

61-2-333-0800

ANIXTER

Anixter Australia Pty Ltd - Sydney, Melbourne, Brisbane, Canberra

Draft DMF Flags Bottlenecks

While the latest draft of Austel's landmark new Decision Making Framework (DMF) was scheduled to be handed down late last month, its first round proposals released in late June suggested significant new definitions of basic carriage service (BCS) and dominance.

The construction of the DMF follows recent amendments to the current *Telecommunications Act* clarifying the operation of dominant carrier pricing rules and giving the industry regulator the power to disallow anti-competitive tariffs.

The first round DMF proposals were outlined at a public meeting in Sydney by consultants engaged by Austel to help formulate the DMF. The consultants include economists Dr Stephen King of Melbourne University and Professor Rod Maddock of Latrobe University, and two Melbourne barristers, including David Shavin, QC.

In essence they recommended that a functional approach to dominance be adopted in place of definitions based around revenue or market share. If incorporated in the Decision Making Framework, their recommendations call for a carrier to be considered dominant where it controls a 'bottleneck' such as an incoming customer connection. At these bottlenecks they would have to file and offer interconnect tariffs at 'flag prices' on a non-discriminatory basis. Carriers will also have to sell themselves their basic carriage services at the filed rate.

Optus Targets 1995 Float

Optus last month named lead managers and a financial advisor for the proposed flotation of the company next year. County Natwest Corporate Finance will act as financial advisors, while the joint lead managers have been named as County Natwest Securities and Potter/Warburg/ABN Amro Hoare Govett.

While some analysts have tipped the first half of next year as most likely for Optus's stock market listing, the carrier would not divulge further details and released a statement insisting that 'All further aspects on the proposed float — including timing — will be determined during the course of 1994, and announced when appropriate.'

At the time of going to press, Optus was also reportedly on the verge of announcing major plans to confront Telecom in the local loop market. Constrained by interconnect charges from competing for local call business via Telecom's network, Optus announced in May an agreement with US cable TV giant Continental Cablevision to explore service provision opportunities. Optus is also known to be looking very closely at wireless local loop technologies.

NetComm Focuses on Core Business

Buoyed by record sales levels in June and a new $3 million UK export order, Australian modem manufacturer NetComm has announced a reorganisation designed to help it compete even more effectively with cut-price US vendors.

MD, Chris Howells, said that NetComm will now be based on two divisions: Domestic Operations, which will concentrate on increasing the company's Australian market share; and Export and New Business, which will focus on the development of export markets and strategic partnership opportunities.

In addition, Howells said NetComm has dropped the distribution of all third party product lines. The decision followed an analysis of all company activities showed they accounted for less than 7% of corporate revenues, contributed little in the way of profit, and had low growth prospects.

PanAmSat's PAS-2 satellite, which was launched by Arianespace on July 8 in Kourou, French Guiana, being given a final check before the installation of its antennas and solar wings. Built by Hughes Aircraft, the 2,920-kilogram Model HS 601 satellite will reach its final orbital location at Longitude 191° West and commence service over the Asia-Pacific region at the end of this month. PAS-2 has 16 63-watt Ku-band and 16 64-watt C-band transponders and can carry over 320 digital channels. It is the region's first private international satellite, providing coverage extending from the western United States through Australia and New Zealand to Asia. It also has a switchable spot-beam for Australia/NZ similar to that of the Optus B1 satellite. The first PAS-2 customers include KDD, Turner Broadcasting System, and Viacom.

Alcatel Wins Brazil Contract

QPSX-developed Metropolitan Area Network (MAN) products manufactured by Alcatel Australia will find their way into Brazil after Alcatel's clinching of a contract with Telebrasilia to partially supply the country's largest high-speed network.

Alcatel Australia will provide MAN switching nodes for the project, which will comprise a 155Mbps SDH transport network, subscriber equipment and network management. Alcatel Australia already exports MAN equipment to Germany, the UK, Belgium, Switzerland, Austria, Italy and Israel.

The network will be implemented in two stages, with access initially limited to public institutions such as government departments and congress.

Troubled Times at Exicom

Exicom issued an explanatory letter to shareholders in late June explaining the company's falling market capitalisation.

The letter revealed that Exicom was looking for 'key personnel' to replace CEO Bob Cruickshank and COO Charles Gilmore. The company is also reported to be up to eight weeks behind on its manufacture of Nortel P-Phones, and had to recall 200,000 of its Touchfone 200 units after they failed in the humidity of Queensland's wet season. In addition, the company was reportedly negotiating the structure of its $30 million debt with several banks.

News of Exicom's troubles forced its share price down from a high of $1.35 in January to around 55c last month.

PacStar to Build Centrex Network

Pacific Star has chosen Telecom to supply equipment for a new Queensland Government telecommunications network which will save the State several millions dollars each year.

The five-year contract, estimated to be worth up to $40 million, involves the supply of a Spectrum Gold voice service and fibre optic network with up to 32,000 lines to 18 different Queensland government departments in Brisbane. The system replaces traditional PABX systems to provide a range of advanced 'intelligent' network services. When complete, it will be one of the 10 largest Centrex networks in the world.

Pacific Star MD, Michael Begun, said the contract could save the Queensland government as much as $12 million in communications costs over the term of the service agreement. The network will use the advanced DMS switch from Nortel, with Spectrum Gold terminals to be supplied by Exicom.

JNA Wins China Contract

JNA Telecommunications has won a contract worth $5 million to supply Digital Data Network equipment to the province of Henan in north central China.

The contract, awarded by the Henan Post and Telecommunications Administration, requires JNA to supply its flagship AS-200 technology platform to support the DDN and form the basis for a Value Added Network.

The DDN will link all the cities in Henan, and connect to China's national network, providing high quality data, voice, fax, video and frame relay services. In the first phase of the project JNA will provide multiplexers, cross connect switches and network terminating units and install some 167 nodes.

Stanilite Electronics has also celebrated a recent success in China, with the announcement of a $6 million contract with the Shanghai Maritime Telecommunications and Navigation Aids Corp to design and supply a maritime safety system.

Local Firms Get ISO Accreditation

A number of Australian manufacturers have recently boosted their international competitiveness by receiving international quality certification for their production facilities.

The research and development arm of Datacraft, Datacraft Technologies, has received accreditation to AS3901/ISO-9001. Officials say it is one of the first in Australia to receive multiple quality standards accreditation, having already won AS3902/ISO9002 certification.

The standards cover the design, manufacture and testing of equipment and software, and the structure and methods of the company's internal operations and customer services.

Melbourne-based communications equipment manufacturer, Design Two Thousand, has also been awarded certification to AS3901/ ISO9001/NZS9001.

Telecommunications vendor Jtec has received AS3902/ISO-9002 certification for its Meadowbank (NSW) factory, covering the manufacture, assembly, testing and rectification of its range of intelligent networking equipment and systems.

ADE Network Technology, a Melbourne-based specialist in internetworking products, has also announced accreditation to AS3902/ISO9002, as has Tasmanian-based Critec, a manufacturer of power conditioning products and services.

John Reidl (left) accepts AS3902 accreditation on behalf of Jtec

In Brief

AT&T has announced it will build a new Asia-Pacific cable manufacturing plant in Queensland. The facility, which will be operational by the first half of 1995 and which will initially employ 40 people, will be the Asia-Pacific region source for the company's Systimax structured cabling products.

Telstra has joined AT&T and Japan's KDD in trials of ATM over trans-Pacific satellite links and submarine fibre optic cable. The trials, which commenced on July 1, will continue throughout 1995 and will involve ATM nodes in Sydney, New Jersey and Tokyo.

Datacraft has signed a technology transfer agreement with US value-added network and equipment provider Network Equipment Technologies. Under the terms of the deal Datacraft will incorporate its Network Terminating Unit technology into NET's Integrated Digital Network Exchange range of products.

Com Tech has automated its entire distribution system and increased overall warehouse capacity by 600%. The new system has increased the speed of order processing by up to 50%, and enables the processing of 40% more orders each day.

Philips has supplied a video link system to Sydney's Long Bay Gaol as part of a pilot to test the use of videoconferencing in court proceedings for bail applications. The system has the potential to minimise the risk of prisoner escapes during transfers.

Acer Computer Australia will establish an Australian facility to design and assemble SIMMS (Single In-line Memory Modules). The SIMMs will be incorporated into Acer PCs sold in Australia, and could also become part of the company's export offerings.

The IPL Group has purchased 60% of Voca Communications from Tubemakers. IPL Datron MD, Stead Denton, said the companies would not merge and would continue to function separately.

Telstra has announced its Switched Digital service (international ISDN) is now available to the Philippines and Ireland, bringing the number of countries customers can access to 19.

Integration Design, a Melbourne-based software company, has won a contract worth around $300,000 to supply and install a computer-based cable management system for 15 RAAF bases and telecommunications sites.

Alcatel TCC has completed laying operations on the PacRim-West cable link connecting Australia and Guam ahead of schedule. The 7,080km of cable was laid in six weeks.

ERG Australia, in conjunction with Fujitsu Australia, has won a contract with the NSW Government to provide a new smartcard system. The companies, who bid for the contract as a consortium called QuickLink, will supply new 'stored value' cards which will include a range of information, from personal identification, medical records and letters of credit, to bus passes for school children.

Telstra Mobile Satellite and Radio Services has signed a three-year contract to supply satellite data services to New Zealand's Ministry of Agriculture and Fisheries. The Satcom-C contract will provide links between 200 fishing vessels and the Ministry, enabling the exchange of information about vessel position, catches of certain fish species, and proposed port calls.

AC&E Sales has been selected as prime vendor in a five-year product sourcing contract to supply Telecom with 75-Ohm 1.6/5.6 Coaxial Connectors and related cords, tools and accessories.

Telstra Enhanced Services will launch a new service for travel wholesalers based on the ATLAS system developed by the Queensland Travel and Tourism Corporation. Telstra has acquired the intellectual property of the system, which has a user base including Greyhound Pioneer, the NT Tourist Commission, the SA Tourism Commission, Tourism Victoria, and Westpac Travel.

Techway has completed a major sale of a rack-mount server system specially designed for areas with limited space to the Royal Australian Navy. The system was developed and built at Techway's Australian facility in Penrith.

The ABC has installed new satellite earth stations in each capital city to reduce the satellite costs of its network. The antenna systems, supplied by Scientific Atlanta, have been engineered for either the Optus Ku band satellites or the Palapa C band satellite, and are capable of both receiving and transmitting.

In Brief

Globalstar, the proposed new low earth orbit satellite system scheduled to provide worldwide communications services to hand-held phones from 1998, has two new recruits — France Telecom and Vodafone. They join nine other companies including Loral, Alcatel, Deutsche Aerospace, Hyundai and Qualcomm.

Alcatel-Alsthom Chairman Pierre Suard has been charged with forgery, corruption and fraud, following an investigation into over-billing of French carrier France Telecom. The company's stock dropped 8.3% on the bourse after the scandal broke.

Northern Telecom has signed an MoU with the New Brunswick Telephone Company (NBTel) to design a broadband multimedia network. NBTel is spending $C300 million on the project.

Arianespace has signed a launch services contract for the new British Aerospace Skynet 4E telecommunications satellite. Scheduled for launch from French Guiana in 1998, the satellite will provide communications services to the British armed forces.

Fiji's new Vodafone-supplied GSM digital mobile phone network was launched last month with a traditional Fijian kava-drinking ceremony. Fiji's President, Ratu Dir Kamisese Mara, made the first call on the system to his daughter in London.

Alcatel has supplied equipment to the city of Antwerp for the first urban ATM network. In the initial phase eight locations are interconnected by two Alcatel ATM switches, providing high-speed voice, data and video communications. Alcatel has also recently won multi-million dollar contracts with MCI and US telco SNET for the supply of SONET cross-connect and other equipment.

Novell has signed an agreement with Ericsson which enables connectivity between Ericsson switches and NetWare servers.

Newbridge Networks has announced results for its fiscal year ending April 30. The company's net earnings for fiscal 1994 were $C157.8 million ($156 million), an increase of more than 160% compared with 1993's figure of $C60 million.

Unisource, an alliance between Dutch, Swiss and Swedish telcos which provides private network services, has joined AT&T's WorldPartners programme. Unisource's European networks will now be linked with the those of WorldPartners.

Motorola has launched a radio-based PCS local loop system in Singapore. The system, called PPS-2000, is said to halve the cost per subscriber of traditional copper-based installations.

Cabletron has reported a 37% increase in sales for its first quarter 1994 over the same period last year. Sales were $US131.5 million for the period ending 31 May, while net income rose to $US36.2 million from $US26.2 million for the same quarter in 1993.

Mitch Kapor's new company, ON Technology, has acquired Da Vinci Systems, known principally for its widely-used e-mail package. Kapor was the original founder of Lotus Development.

Nokia's interim financial report covering the period from January to April 1994 shows net sales growth of 33% to FIM8,565 million ($2,314 million), attributable mainly to strong growth in the company's Telecommunications and Mobile Phones divisions.

MEASAT-1, Malaysia's first telecommunications satellite, is scheduled to be launched by Arianespace at the end of 1995.

European consortium Orion Atlantic has launched what it claims is the first satellite-based ISDN switched service. Orion Atlantic Express is a bandwidth-on-demand service linking private networks with any ISDN public network in Europe.

Eunetcom, the joint venture between Germany's DBT and France Telecom, has won its first major contract, a $US200 million deal to manage Dun and Bradstreet's data communications. Under the five-year arrangement, Eunetcom will manage D&B's existing network and provide services such as X.25, managed links, frame relay, router management, EDI and e-mail.

Olivetti has joined with Hughes Network Systems to form a joint venture company to provide VSAT satellite communications services. The first offering from the new Hughes Olivetti Telecom will be Hotstar, a platform for private voice, data and video.

Telecom Finland has decided to close down its Pointer CT-2 communications network because of a lack of subscribers.

Bell Atlantic, Nynex Join Forces

Hot on the heels of its aborted merger with cable giant Tele-Communications Inc., Bell Atlantic announced it will combine its cellular services properties with Nynex Corporation to form a US nationwide wireless communications company.

The new entity, if approved by US regulators, would be second in size only to McCaw Cellular Communications. The combined market share of both companies already covers a population of over 55 million, and the two have jointly operated a successful cellular system in metropolitan New York, the largest single cellular market in the US.

The new company will adopt an aggressive strategy for acquiring and operating personal communications services (PCS) licences, which are expected to auctioned by the Federal Communications Commission later this year. It then plans to coordinate newly won licences with existing cellular operations to quickly establish a coast-to-coast PCS network.

Together, Nynex and Bell Atlantic should have the financial and technological clout needed to compete in this highly competitive market against formidable rivals AT&T-McCaw and MCI-BT-Nextel.

The cellular merger is the first time since the divestiture of AT&T 10 years ago that two of the US Regional Bell Operating Companies have attempted to form such a large scale alliance.

BT and MCI Launch Concert

The alliance between BT and MCI has already borne fruit in the form of a $US1 billion joint venture called Concert, which will offer a full portfolio of global communications services for multi-national customers.

Concert's network will link more than 5,000 network access points in 55 countries by mid-1995, and Global Customer Support Centres have already been established in London, Paris, Tokyo, Sydney and Cary, North Carolina.

Concert's first offering, the Virtual Network Service, was launched in Australia, the UK, France, Belgium, Germany, the Netherlands and Sweden early last month. It includes advanced corporate voice services, customer-defined call handling, and corporate calling cards with operator services, and audio- and videoconferencing services will be offered in the future. The network will also offer around-the-clock multilingual customer support and a global integrated billing and a reporting system tailored to customer requirements.

The network's first customers are expected to be Holiday Inn, First Data Corporation, Glaxo and British Petroleum.

NT, Ericsson Triumph in Taiwan

Canada's Northern Telecom has won its largest single GSM telephone system contract so far. Announced by Taiwan's Long Distance Telecommunications Administration, the contract for the 500,000-subscriber system is estimated to be worth more than $US100 million.

The new system will provide GSM cellular service for all the major urban areas of Taiwan, and is expected to become operational in the first half of 1995. NT will supply a turn-key solution including radios, switching equipment and associated hardware, as well as engineering, installation, commissioning and maintenance.

Ericsson has also secured a large contract for the supply of a digital EDACS trunking radio system to the Taiwan National Police Administration. Worth $US110 million, the project will have island-wide coverage, and includes eight multisite subsystems, over 80 repeater sites and over 40,000 radios.

AN ADDITIONAL TWO BILLION PEOPLE COVERED BY THE FREE ENTERPRISE FOLKS WHO BROUGHT YOU PAS-1

PanAmSat

is proud to announce
the successful launch of its
Pacific Ocean Region satellite
on July 8, 1994.

PAS-2

STATE-OF-ART
COMMUNICATIONS SERVICES.
THE MOST POWERFUL
INTERNATIONAL SATELLITE
IN THE PACIFIC.

PANAMSAT, L.P. • ONE PICKWICK PLAZA • GREENWICH, CT. 06830 • USA
1/203/622/6664 • FAX 1/203/622/9163

Jtec

Solutions

What Keeps BP Australia On The Move?

Jtec's communications network is BP's efficient business engine.

BP Australia's commitment to the highest standards of quality in its products, service and distribution, keeps the company on the move throughout the 90's and is earning it an enviable reputation.

Behind much of the company's success lies its sophisticated communications network, an intricate highway that links its operations across the country and around the globe. In this area, as with its products and services, BP Australia demand the best.

When it came to choosing a communications partner with an understanding of its business issues, the choice was simple—

Jtec.

BP Australia relies on its communications network. It is the key to maintaining an efficient business network. The key to maintaining BP's winning edge. That's why BP relies on Jtec.

Jtec. Keeping BP Australia On The Move.

Head Office
Jtec Pty Limited
ACN 003 169 088
Unit 3
118 - 122 Bowden St
Meadowbank 2114
NSW Australia

Phone

Sydney	61-2-809-6933
Melbourne	3-263-0200
Brisbane	7-229-9233
Canberra	6-253-3399
U.K.	44-494-473757
Europe	353-21-277422

Fax

Sydney	61-2-809-6619
Melbourne	3-263-0215
Brisbane	7-363-0023
Canberra	6-201-3003
U.K.	44-494-536254
Europe	353-21-270420

Privatisation Back on the Agenda

Despite opposition from the CWU and some elements of the ALP, the privatisation of Telecom may soon become government policy.

Is privatisation of Telecom on the ALP agenda or not? It's hard to tell these days, but the difficulty of answering that question means that it is certainly on the media's agenda.

Don Greenlees in *The Weekend Australian* of June 18-19 was unequivocal. He opened a major article by saying: 'Forget the denials from the Prime Minister, Paul Keating, this week that Telecom is not for sale. The battle within the ALP to save Telecom from the privatisation brigade is already on and will be fought all the way to the ALP National Conference in September.'

He may well have been right. Less than a week later Keating was, to say the least, equivocal on the issue. In an interview on ABC *Lateline* on June 22 he was asked by Kerry O'Brien: 'How will Telecom fund its expansion overseas which is, presumably the only way it can maintain its pace among the world's major telecommunications carriers?' Keating said the ownership of Telecom was "not especially important" in the post duopoly era. What was important, he said, was that Telecom was "subject to the competitive disciplines of that particular marketplace . . . Telecom is finding it can fund itself. It's going in and it's making innovative deals around the place with various private sector organisations."

O'Brien persisted: "So it doesn't matter if it's publicly owned or privately owned?" Keating replied: "Not of its essence, no." And so it went on, with the PM never actually saying no to privatisation but trying to steer the question to which bits might be sold off. These words contrasted with responses given to Shadow Communications Minister Richard Alston by Senator Gareth Evans during Question Time in the Senate on 20 June. Alston asked: "Does the Minister give an unequivocal commitment that all parts of Telecom will remain wholly government-owned for as long as Mr Keating is Prime Minister?" Evans replied: "That is the Government's position. It has been made perfectly clear by the Prime Minister and the Minister for Communications and the Arts and it is made as clear as necessary by me."

He must have been fuming after Keating's June 22 performance and so must have Brian Howe, Deputy Prime Minister and leading light in the ALP Left. Apparently unwilling to let the Prime Minister make all the running, in mid-July he declared in a much-publicised speech that "privatisation is not the answer for Australia in the 1990s."

It was also reported in the middle of last month that Communications Minister Lee had drafted a policy position for the ALP's September conference which substantially strengthened the party's commitment to keeping Telecom in public hands.

Keating's remarks also contrast starkly with the election campaign-inspired commitments he gave in a letter to the Communications Workers' Union in February 1993 when he said: "Telecom will not be sold under my Government. Indeed we are committed to its growth and success in the public sector. The Government's investment of $2 billion in Telecom is proof that we want to see a strong, commercially viable publicly-owned company. Privatisation will only be a distraction at a time when Telecom needs to focus its energy on the new competitive regime."

Death of a Thousand Cuts

If nothing else, the *Lateline* performance was a handwashing of any Government responsibility for the future of Telecom which, if it remains wholly government-owned, will have to operate in an increasingly competitive and very capital intensive marketplace with only two sources of funding: revenues and loans. It will have no freedom to seek additional equity as and when its directors feel the time is ripe, unlike Optus or any other public company.

As one commentator on the privatisation debate remarked: "There are worse fates than privatisation." Yes, and death by capital starvation would be one of them.

Death by asset stripping would be another, and there are many who believe that this is the sort of privatisation on Keating's agenda. They suggest that some of Telecom's so called non-core assets could be sold off, like the mobile business and its directories, the two most profitable and easily separable businesses.

Who says they are non-core? Everyone knows that the world is moving towards personal communications, where a telephone number is associated with a person, not with a box stuck on the end of a wire. Any carrier without a mobile business will have a reduced role in such a future.

Richard Alston blamed Keating's attitude to Telecom on a grudge dating from his

Stock Watch

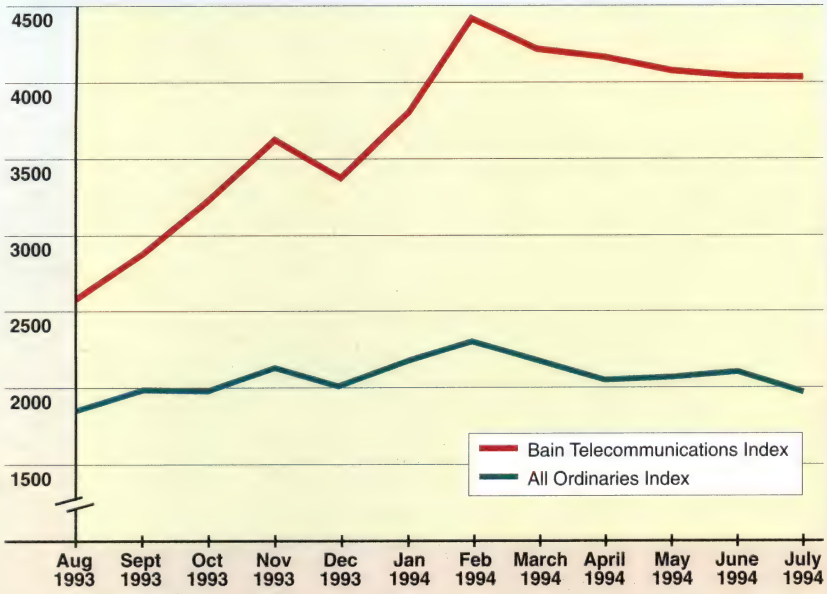

The Bain Telecommunications Index (BTI) includes only those public companies which derive the majority of their earnings from the telecommunications sector. It is calculated in the same fashion as the All Ordinaries Index and is maintained by the Australian Stock Exchange Index Office. The formula is:

$$\text{Today's Closing BTI} = \text{Yesterday's Closing BTI} \times \frac{\text{End of Day AMV}}{\text{Start of Day AMV}}$$

where AMV refers to Aggregate Market Value - i.e. total market capitalisation. The start of day AMV is the previous day's end of day AMV and is adjusted for any changes in the BTI such as additions, capital reconstructions, rights issues etc. This ensures that any movements in the BTI are the result of trading activity in that day only.

Company	Market Capitalisation $M	Performance %	Share Price ($) 1/7/94	52 Week High	52 Week Low
AWA	163	19.4	0.80	1.02	0.62
Datacraft	112	320	1.05	1.13	0.23
ERG Australia	438	90.1	3.38	3.70	1.54
Exicom	65	-4.4	0.65	1.40	0.50
JNA Telecommunications	68	9.5	2.30	3.40	2.08
Matrix Telecommunications	123	112.5	2.55	3.65	1.20
NetComm	6	N/A	0.68	1.40	0.55
Scitec	24	328.6	0.30	0.40	0.07
Stanilite Pacific	199	87.7	6.15	6.70	3.32
Techniche	92	164.5	4.76	6.50	1.84

Market capitalisation is the share price multiplied by total ordinary shares on issue.
Performance is the percentage share price movement over the past 12 months.

Source: Bain & Company

unsuccessful attempt to keep Telecom and OTC separate in 1990.

But it seems unlikely that Keating's motives are so petty or so unsubtle. Keating and the 'privatisation brigade' in the ALP are more likely to be deliberately spreading rumour and speculation about piecemeal privatisation to sugar the pill of wholesale privatisation to make it more palatable to those who oppose it.

And they are succeeding. *The Age's* left-wing economic commentator Ken Davidson seems to have swung in favour of privatisation of the enterprise as a whole as the only alternative to a carve up. On June 18 he wrote: 'The changes taking place globally will open up opportunities for world-class telecommunication companies . . . Telecom either gets into these markets now or it loses out. The pre-condition for such a move is privatisation, so Telecom is free to take commercial risks and raise capital for expansion, subject to market judgments and disciplines. Even more importantly, the existence of shareholders, who have rights to protect, gives some guarantee that the interests of Telecom (which in most cases coincide with the public interest) will be taken into account when the new local regulatory regime is established in 1997 after the present review. The alternative is to asset-strip Telecom by selling off the profitable bits and leaving the husk in the form of an unprofitable fixed local network. Then, like the railways, it would need constant cash injections from government or higher local charges to cover the estimated $300 million annual losses on the local network.'

Joining him in this change of heart is Kevin Morgan, long-time advocate of public ownership and for 10 years (until his resignation last November) an executive of the CWU and its predecessors (see *Opinion* on page 61).

So has the trick worked? Is the Communications Workers' Union buying into wholesale privatisation and manoeuvring to get the best possible deal for its members and its ideology (i.e. a generous employee share scheme, employee representation on the board and strict limits on foreign ownership)? Not everyone thinks so.

The recurring theme among those opposed to privatisation piecemeal is that it would reduce Telecom's ability to be a significant force overseas, and thus a significant supporter of Australian exports of high technology and telecoms services. But to advocate wholly public Australian ownership as the key to success is to underestimate the scale of the forces shaping global telecommunications. A strategic alliance with a large overseas carrier taking a significant stake may be the best way to secure that future.

One great difficulty facing the government will be an accurate valuation. Regulation has a great impact on the value of a telecommunications carrier and Australia's regulation is at present in a great state of uncertainty. Until the review into the post-1997 regulatory regime is completed Telecom's valuation will be a moving target, kicked along also by the great uncertainty surrounding its big buck investments in information superhighways, Pay TV, video-on-demand etc.

Kim Beazley, minister responsible for asset sales, used exactly this argument to re-assure the ALP that there would be no privatisation before 1997. Quoted in the *Sydney Morning Herald* he said: 'If you want to encourage the competitive environment you can't mix it with privatisation of the main carrier . . . that's the mistake the British made.'

Stuart Corner is the Managing Editor of Exchange and Editor of Telenews Asia.

Telstra Staff Levels

Consultation Replaces Cuts

Recently stung by a survey which showed Telstra's performance ranked poorly against other leading international telephone companies, Telstra management may find some comfort in the fact that the survey is to be re-worked, and its statistical base broadened. The likely result is that Telstra will emerge looking somewhat more efficient than it did in the first survey.

Among the original findings of a report by researchers Strategic Information was that Telstra ranked 13th out of 15 in the number of telephone access lines per carrier employee. Top of the performance list was Ameritech, followed by Bell Atlantic with around 250 lines per employee. Telstra, with 66,000 workers and around 125 lines per employee was ahead only of New Zealand Telecom and Malaytel. NZ Telecom's ranking should be improved by plans to cut staff numbers down to 7,500 by 1997 from the present 9,000 and its 1987 peak of 25,500. However, Telstra has recently come to a loose arrangement with the Communications Workers' Union to maintain present staff levels for the next year or so, meaning its access line/employee rating will also remain static, or even decline.

Clearly aware of Telstra's displeasure at his survey results, Strategic Information Director, Nigel Fitzpatrick said the company now planned to extend its research to between 35 and 40 telcos, rather than the original 15, in a final report to be published at the end of August. Fitzpatrick said a major problem with analysing Telstra was that it did not publish as many figures in its annual report as other international telcos. "For example, they don't publish access lines per operating employees as the others do. They don't even have levels of employment in their annual report, so we had to ask relatively senior section heads within Telstra to establish that figure."

Telstra's Chief Operating Officer, Consumer and Commercial, Charlie Zoi, told *Australian Communications* the Strategic Information survey was like comparing apples and oranges. "Bell Atlantic is not comparable to Telecom," he said. "They don't have long distance services like IDD and STD, whereas we have the lot, including mobiles, domestic and offshore business. It might have been easier in the past to make these comparisons, but no two companies are the same anymore. Some would be very labour intensive because they don't contract work out, for example."

Zoi said customer demand was determining the size of the Telecom work force, and this was evident in Western Australia, Queensland and country NSW where business was booming. "What we are doing in each unit is measuring the workload and making forecasts of the staff levels needed. The workload has dropped in areas like network maintenance, but risen in the sales and customer service areas. We're spending a lot of money training people, building new plant, replacing the customer access network, starting up the Visionstream operation, basically building the company up. Our export business should be generating around a billion dollars by the end of the century and so we need more people in the International division. Annual revenues of $13 billion from 66,000 employees is a pretty good indication of our productivity levels — close to $200,000 per person."

Federal Secretary of the CWU, Paul Watson, confirmed that Telstra would retain its present staff levels for the moment, but this was not the result of any direct negotiation, more a recognition that if service and quality were to be delivered, staff needed to be on hand to deliver it. "They were setting budgets and staff levels from a strictly monetary point of view but we convinced them this should change if they were fair dinkum about service and quality. We did a bottom-up staffing model with Consumer and Commercial and that had a fair role in stabilising the work force."

Watson said other contributing factors in maintaining staff levels were the appointment of former Liberal Minister, Ian McPhee, as Human Resources Director as well as the work of the recently-formed Jobs Task Force, with representatives from Telstra, the CWU and the Community and Public Sector Union. "McPhee has an consultative approach to management, which has been lacking for some time. We can already notice a change of attitude," Watson said. "The Task Force will look at how to handle any future dislocations or new growth areas to see if staff can be re-directed from within Telecom or from outside. They have a fairly wide brief and they'll also look at potential non-core business areas which Telstra could develop on a sound commercial footing." The Task Force will report to Telstra management in about three months.

Despite the Strategic Information findings, indications of productivity improvements at Telecom emerged at an International Telecommunications Society meeting in Sydney last month. The Bureau of Transport and Communications Economics noted that while total factor productivity (TFP) estimates for Telstra were not available for more recent years, some partial productivity measures indicate significant increases in productivity since the 1991 telecommunications reforms.

According to Research Director in the Bureau's Communications Branch, Kym Starr: "The number of telephone services per average number of employees for Tele-

com increased at a compound annual rate of 10.2% between 1990-91 and 1992-93, compared with an annual rate of 6.1% between 1987-88 and 1990-91. The average number of employees at Telecom fell at a compound annual rate of 6.5% between 1990-91 and 1992-93, compared with an annual fall of 1.8% over the earlier period."

Just as Telstra's staff numbers have fallen from a mid-1980s peak of over 90,000, total employment in the Australian communications industry sector also appears to be gradually declining. Australian Bureau of Statistics figures, based on the monthly Labour Force Survey which covers 33,000 private and non-private dwellings around the country, show total communications sector employment has dropped from 146,000 in the fourth quarter of 1990 to 130,000 in the fourth quarter of 1993. By the second quarter of 1994, the figure had dropped to 120,000.

A different ABS survey, based on a sampling of 13,500 private and public sector communications employers, shows employment has dropped from 128,800 to 106,000 over the same period. According to Starr, the discrepancy arises from the fact that the Labor Force Survey treats communications as one small part of a much larger picture, and the survey of major employers provides a more accurate indication of trends in the telecommunications industry.

"The communications sector is far bigger than just the carriers and includes downstream service providers such as Australia Post, the radio communications industry, even broadcasters," he said. "Among the carriers there has been shrinkage and even the 3,000-odd Optus employees and the handful at Vodafone don't make up for the cuts at Telstra. In the wider communications sector, there has been fluctuating levels of growth driven by new services, but the overall employment trend is down."

Bernard Levy

New Zealand

BellSouth's Burnt Fingers

BellSouth executives must rue the day they heard of New Zealand. Thanks to a mixture of incompetence, competition and downright bad luck, the US telco has not got off to a good start in that country.

BellSouth's first error was to bid for the wrong frequency. The company tendered for, and won, TACS-A. At the time it could have had TACS-B. That would have been smarter. In places where AMPS and TACS cellular networks co-exist, there is interference. Avoiding this requires setting aside some bandwidth. Telecom Corporation of New Zealand (TCNZ) was the incumbent on AMPS, so the onus to trim bandwidth fell on Bellsouth. This decreased the maximum amount of traffic BellSouth could sustain on its network.

At the time, few knew of the potential interference. Yet, you might expect one of the world's largest telcos to be aware of the problem. What makes this more disturbing is that BellSouth Enterprises, the company's overseas investment arm, specialises in cellular networks. In addition to New Zealand, it has investments in the US, South America, France and Denmark.

Eventually the non-intervening New Zealand Government intervened and switched BellSouth to the TACS-C band. Within a week, BellSouth announced a major investment programme. Did BellSouth do a secret deal with a 'hands-off' government? Industry insiders think so. BellSouth certainly made noises about pulling out of New Zealand when the interconnect agreement with TCNZ reached a difficult stage. On the other hand, BellSouth spokesman James Norman said the band allocation broke the dam holding back further investment.

If BellSouth's engineering erred, so did its accountants. When it started operations, executives told journalists BellSouth could build a national cellular network for between $NZ50 and $100 million. In 1992, company marketing blurb mentioned spending some $NZ200 million. Now, BellSouth says it will have invested some $NZ300 million in New Zealand by March 1995.

Jamie MacDonald, General Manager, Marketing for TCNZ's Telecom Mobile said, "New Zealand's mobile phone customers are a demanding lot." He admits his company misread the market in 1989. Then, TCNZ assumed mobile telephony meant car-phones. "In fact," he said, "New Zealand is largely a pocket phone market."

Patchy Coverage

BellSouth's network is growing rapidly, but for some unexplained reason, the company seems to have had more trouble getting planning consent for its sites. There was even one case where a medium complained that a cell site might disturb his aura. On the whole, these problems now seem to be under control.

In February, a weekly newspaper, *The Independent*, published a map showing BellSouth's Auckland coverage was still patchy some eight months after starting operations. The associated story pointed out that the high profile Wellington launch had been put back from November to February and then to April. Just days after the story appeared, BellSouth issued a press release saying the Wellington service had started. However, by mid-July there was still no high-profile Wellington campaign. Rumour has it that BellSouth is still experiencing technical difficulties offering adequate coverage in the hilly harbour capital.

BellSouth says it will offer coverage in 24 of New Zealand's largest towns and tourist destinations by March 1995. The latest expansion phase will cost $NZ50 million and follows the $NZ30 million earmarked for Christchurch. Work there is already under way. Over the past 18 months, a further $NZ200 million has been spent setting up networks in Auckland, Wellington and Hamilton. By early next year its network will cover 75% of the New Zealand population.

But critics have reserved their harshest words for the company's high profile marketing. When BellSouth first started marketing its GSM cellular service last winter, New Zealanders were promised "tomorrow's technology today." The company spent some $NZ3 million promoting its message. In New Zealand terms, that's a huge budget. During the weeks leading up to the launch, Auckland billboards proclaimed "tomorrow is 11 July, 1993." After that date that message switched to "this is tomorrow." TV advertising used the Roxy Music song of the same name. It was a glossy and expensive attempt to sell a digital cellular service.

Conventional wisdom says this should work in New Zealand. This is, after all, a nation of early adopters. On the surface at least, the campaign flopped. Some flaky sales techniques by BellSouth's external dealers didn't help. One newspaper reported salesmen telling customers that GSM phones could automatically translate Japanese to English.

In February 1994, after eight months of operation, BellSouth had less than 2,000 subscribers. At that time, Telecom Mobile boasted some 135,000 connections. In marketing terms, it cost BellSouth $NZ1,500 to 'buy' each customer. At the time of writing Telecom Mobile had 160,000 subscribers. BellSouth won't disclose customer numbers, but outsiders estimate there are well under 5,000. As one rival cellular phone retailer unkindly put it, "you're more likely to see Elvis than a business card with (BellSouth's) 021 prefix."

That's rather unfair. BellSouth is still very much in its start-up phase. It can hardly attract large numbers of users until its coverage improves. With the allocation of TACS-C that's now more likely to happen. The past might be unsatisfactory, the investment may have been higher than expected, early sales may be slower than desired, yet BellSouth still has everything to play for. Both Telecom and BellSouth estimate New Zealand will have 500,000 cellular connections by the end of the decade — that's 340,000 more than today. Even a slender slice of that pie will be a valuable prize.

Bill Bennett

Step up to Fibronics

Ethernet, Token Ring, FDDI and ATM Solutions

- GigaHUB
- MultiHUB
- Token Ring SNMP Stackable CAU
- Ethernet 10BaseT Stackable Hub
- FDDI Workstation Server
- InterView Network Management
- ...and more

Available from qualified Network Integrators.

Now distributes and supports Fibronics in Australia Call Olivetti Network Integration on (02) 368 1988 for product information. We can recommend a Network Integrator to customise your Fibronics solution.

Resale

AAPT Revenues Rocket

Leading value added service provider, AAP Telecommunications, has scored impressive revenue increases over the past three years and is forecasting more of the same in the lead-up to greater competition in 1997. With new switches and infrastructure installed, an aggressive telemarketing campaign, and lower average prices between major interstate locations than Telecom or Optus, the company has already established an edge over potential post-duopoly rivals.

According to the company's Chief Executive, Barry Wheeler, revenues grew by 120% in 1992 and 140% in 1993. A similar figure was expected for 1994-95, though such a steep curve would have to plateau off at some stage, he said. With a customer base of 6,000, ranging from hundreds of small businesses spending around $400-$500 per month on long distance calls, to 80-odd large corporations and government departments spending up to $400,000 per month, Wheeler said the company was now trading profitably.

Wheeler said AAPT had almost doubled its revenues last year to about $42 million, though investment and operating costs had contributed to a loss of $5.6 million. This compared with a loss of $8.5 million in 1992, when revenues were $23 million.

Wheeler attributed the projected 1994-95 profitability to the installation of DSC 600E megahub switches in Sydney, Melbourne, Brisbane and Perth, with upgrades also scheduled for Adelaide and Canberra. Built to accommodate high traffic volumes, the 600Es have replaced older, less powerful units from Nortel, which will be relocated in regional centres such as Newcastle, Wollongong and Ballarat by the first quarter of next year. Wheeler said that as a result of the switch upgrade, the company would probably put on about another 50 staff in sales, technical support and customer service, bringing the total to around 300, from the original 20 when the company started up in 1991.

"The really big increases in revenues did not come until after 1992 when we moved away from dedicated leased access lines to our switched services such as FirstChoice, Vista, OneGovernment, Assurance and our virtual private network service, Vantage, with the 1414 access code," he said. "The 600Es are carrier switches used by MCI in the United States. We bought big capacity and they're compatible with the network intelligence and the MCI data access points we use."

Much of the funding for the switch upgrade came from an injection of almost $11 million from AAPT's shareholders, including AAP Information Services with 51%, and MCI and New Zealand's Todd Corporation both with 24.5%. MCI and Todd both have stakes in New Zealand's second carrier, Clear Communications.

Wheeler said AAPT wanted Telecom's bottleneck control of the local loops removed, and the price of access to the loop made equal for all comers. "Similarly on the technical side, we should be allowed CCS #7 interconnect. If we had preselection ballot functionality, we'd be able to drop the 1414 access number, which requires that we invest in auto-diallers for our clients or put a software fix into their PABXes. We carry this cost, but it's one we could do without."

Rejecting any suggestion that the present partial deregulation had created a 'cosy' environment for Telecom, Optus and AAPT, Wheeler said his position was 'the more competition, the better.' "A lot more deregulation has to happen, and if you only have a smaller number of players, the pace will slow down. No-one should imagine the business is neatly sliced, diced and apportioned. It isn't. Even the major corporates and government departments, which are our biggest revenue generators, are learning to play the competition game very well and spreading their business across all the players, which is smart."

Bernard Levy

AARNet

Too Popular For its Own Good?

Control of the Australian Academic Research Network (AARNet) may shift away from academics to commercial users as a result of the explosion in demand for access to the Internet by non-academics. This and other issues such as tariffs and charges are certain to be hotly debated at a meeting planned for Canberra this month involving AARNet's owners, the Australian Vice Chancellors Committee (AVCC), the carriers, equipment suppliers and large and small users, under the aegis of the Australian Science and Technology's (ASTEC) Working Party on Research Data Networks (RDNs).

A distinct possibility also exists that other means of accessing the Internet will be established independent of AARNet, either by Telstra Enhanced Services or by commercial interests seeking to provide a highly reliable service for the commercial and non-academic sector. Adelaide-based Ausnet Services, with the backing of Hong Kong financiers, is conducting a feasibility study to establish whether such an alternative service would be viable and whether it would buy bandwidth from AARNet or go directly to Telecom or Optus.

How AARNet develops in the short to medium term hinges largely on the agenda and future directions established by the ASTEC Working Party and its Chair, Professor Anne Henderson-Sellers of Macquarie University, who have conducted an exhaustive study on RDNs since last November. The study's terms of reference were to examine Australia's future requirements for RDNs, including their role and use in universities, government research organisations, education, industry and the wider community. The critical issues of whether to charge, how to charge and who to charge have also been canvassed but will remain in a state of flux for some time yet. The Working Party also hired consultants, Cutler & Company, to identify the issues involved (see 'Building Australia's R&D Network' starting on page 75).

While Professor Henderson-Sellers had originally supported the notion of volume charging on AARNet in a draft report released in April, so many protests were registered that she changed tack and will now try to convince the Federal Government to continue some form of financial support beyond the scheduled cut-off on January 1, next year. "Once the draft was released, the number of responses was more than double the number of submissions, which was quite extraordinary. Most were against volume charging, but I understand the AVCC still plans to go this way," she said.

"I can see both sides of the volume-charging argument, but my fear is that volume charging would tend to curtail research discovery, which is what networks like the Internet are all about. There is a middle ground and my view is that charging should be geared to the level of urgency of the information you want to receive or transmit. If the network is not congested, you don't pay. If it is congested, you can be offered the option of paying to ensure that a real time transmission continues uninterrupted. Those transmission costs could be covered by research grants, etc., and people with little money, such as students, could still use the network for free."

However, Henderson-Sellers conceded she may have a tough job convincing the AVCC and the AARNet Board at this month's Canberra meeting. She said her preferred option was for a national consortium to administer AARNet, with the AVCC still taking a major role, but accepting input from the major users and carriers.

The ASTEC Working Party's final report to the Prime Minister's Office is expected by September. Other reports on telecommun-

ications infrastructure needs and issues, including those from the Broadband Services Export Group, the Bulletin Board Task Force and a Senate Standing Committee are also due to be delivered later in the year.

Treasurer of the Australian Unix and Open Systems User Group (AUUG), Frank Crawford, said his group wanted to see the AVCC relinquish control of AARNet to its major user groups. This would facilitate more widespread use of the Internet in society generally, especially in multimedia communications. "While AARNet is operated by the AVCC, it will always be regarded with some suspicion by industry as a toy for academics," he said.

AARNet's General Manager, Peter Saalmans, estimates that AARNet has up to 500,000 users, connected through 100,000 host systems from PCs to supercomputers. Traffic was increasing by around 130% per year, and doubling every eight to nine months, he said. Describing AARNet as a "national asset with economic implications that cannot be ignored," Saalmans said the network was constantly being upgraded to provide more bandwidth. AARNet's plan to boost capacity by installing T3-type 35-50Mbps microwave links at university campuses and research establishments was a significant step. "Telecom doesn't have a Fastpac product above 10Mbps at this stage, whereas 35-50Mbps microwave products already exist. So apart from partially freeing us up from the carriers, microwave provides a high-powered solution for more complex applications."

While Telecom has been sent a Request for Information by AARNet, as have Siemens, Ericsson, AAP Telecommunications, ATI and others, Telecom appears less than enthusiastic about the microwave option.

AVCC: It's Ours

The present debate is very much about balancing the future research networking needs of Australia as a whole with the AVCC's desire to remain master of AARNet's destiny. The AVCC's position is hardly surprising, given that it built the network to serve its own community. Something of a power struggle between the AVCC and Telecom emerged about two years ago, when the Government made $13 million available in the Budget to upgrade AARNet and support Research Data Network Cooperative Research Centres around the country.

While Telecom immediately advanced its metropolitan area network product, Fastpac, as a total network solution, the AVCC refused to be bullied, demanding the right to choose from all available technologies. In the event, as part of a $7 million upgrade over the next 18 months, the AVCC has committed to spending over $3 million on 2Mbps and 10Mbps Fastpac solutions for lines linking the State capitals with the national hub in Melbourne.

Professor Don McNicol, President of the AVCC and Vice Chancellor of Sydney University, said the rapid growth in the number of non-academic users would inevitably place pressure on AARNet, and alternative access routes to the Internet may be the solution. "As soon as we lease more Internet bandwidth, it is immediately saturated. It's quite possible that AARNet could remain a purely academic network sitting beside other networks," he said. "But different networks require gateways between them and from the research community's point of view, they want to stick with one network as much as possible."

McNicol conceded that because of the growing proportion of non-academic users, there may need to be changes to the AARNet Board, taking in government and private sector representatives.

On the question of volume charging, McNicol said he hoped this would be introduced as planned at the beginning of January next year. "If you analyse the present arrangements, we are not fairly charging

the various members. For example, Sydney University does better at the moment than if it was volume-charged. From AARNet's point of view, the important thing about volume charging is that it is institutions, rather than individuals, who are being charged. Individual charging is enormously difficult and what you need is a low-cost charging system that optimises the usage of the network with a combination of off-peak and peak rates."

AARNet presently supplies bandwidth to around 10 resellers, including Connect.com, OzEmail, Pegasus, Dialix, Australia On Line and others at the retail end of the market. Ray Gatt, Managing Director of potential rival Internet wholesaler, Ausnet, said his biggest concern was pricing — particularly AARNet's intention to now go ahead with volume charging next year.

Managing Director of OzEmail, Sean Howard, said he supported greater competition in the provision of access to Internet. "I'm not sure AARNet should be the sole national access point. It might be good for Telecom to get more involved, just as British Telecom has done. But we shouldn't be carried away about who lays the wires."

According to Connect.com's Managing Director, Hugh Irvine, there's no such thing as a free network. "The first colloquial myth floating around is that Internet should be free. In my opinion, network usage should fund the network itself. But there are no simple answers."

Deputy Secretary of the AVCC, John Mullarvey, said there was no question more bandwidth was needed to provide optimal Internet access for Australia. "With any network, someone has to pay for it and up till now, we haven't been able to find anyone willing to do so. If Professor Henderson-Sellers can convince the Government to put up the money, good luck to her. We haven't been able to do so and we won't hand AARNet to anyone else while we are paying for it."

Bernard Levy

Industry Development

Local Companies Join FTA Club

The more pragmatic approach to local industry development recently adopted by the Australian Government represents a concession that multinational companies will accept only so much government direction on how they should conduct their commercial activities.

By now allowing Australian companies, as well as the multinationals, to join the Fixed Term Arrangements (FTAs) and achieve 'endorsed supplier' status, the Government has signalled a renewed willingness to generate as much industry development as possible from the purchasing policies and strategies established over the past half-decade. It's not insignificant that the decision to allow companies to migrate from the somewhat rigid Partnerships for Development Programme (PDP) and the Industry Development Arrangements (IDAs) to the FTAs, has been taken under the Ministerial reign of Senator Peter Cook at the Department of Industry, Science and Technology (DIST). The Senator was closely involved in the GATT negotiations, and his department has been assessing whether Australia's industry development arrangements and offsets schemes may be inconsistent with GATT.

According to Senator Cook's IT adviser, Glen Carlos, while the new arrangements might been seen as encouraging imports of cheap IT products, the benefits to Australia outweighed such concerns.

"The PDP/FTAs caused companies to set up assembly and R&D in digital cellular phones in Australia and we couldn't have achieved that if we hadn't push them down the path of strategic investment. Naturally, we'd love to have all the assembly work done here, but we have to be realistic and we'd prefer to become the site for the smart end of the business. It's not that we aren't concerned about cheap imports, we're more concerned about strategic investments."

Carlos said endorsed supplier status for Australian companies would have an impact not only overseas, but in Australia itself. "We found documented cases where State Governments were effectively excluding Australian companies because they required them to be involved in Partnership-type schemes and they couldn't get in."

Under the PDP, whose imminent death has been rumoured for some years, vendors to government have been required to meet rigid export and R&D targets if they sold between $10-$40 million to Government annually. The IDAs required commitments to local industry development in exchange for being allowed to import customer premises equipment (CPE) such as mobiles, PABXs and small business systems.

Under the FTAs, started in 1991, companies can nominate their own programme of investment in Australia on a 'best efforts' basis, with favour bestowed on plans for long-term investments in the country.

Director of Telecommunications and Broadcasting at the Department of Science, Raph Cicchini, said half a dozen of the 19 companies involved in IDA scheme had migrated to FTAs and others were expected to follow soon.

AT&T has recently signed up, committing itself to a four-year plan involving more than $104 million in industry investment and development. The AT&T plan has three planks — a $20 million manufacturing plant for its advanced cabling system, Systimax; a new research and development centre at the company's Global Business Communications Systems premises in Sydney aimed at enhancing the Definity multimedia switching system; and a commitment to partnerships with local companies and setting of an export target of $74 million in its export drive into the Asia-Pacific.

Among others to join the FTA scheme are Lexmark, which will locate its Asia-Pacific headquarters in Australia and build printers here, and Novell, which will establish a regional headquarters and testing centre.

Australia companies to join up are IPL Datron, which distributes Oki printers, fax machines and mobile phones; Kambrook, which produces household and electronic goods; MHG Plastic Industries; and Cellular Services & Distribution. Other local companies are known to be queuing up to join. In Kambrook's case, the company earns the right to import cellular mobile phones by manufacturing telephones in Melbourne which it ships to Bangladesh. The company's Corporate Strategy and Business Development Manager, Barry Davies, said having access to CMT imports would provide cash flows which would help grow the company's information technology business and fund its entry to new areas such as gaming machines.

The viability of the IDA/PDP/FTAs is being reviewed in the light of Senator Cook's White Paper on Government Purchasing Initiatives. The review is being conducted by the Bureau of Industry Economics (BIE) in consultation with the Department of Communications and the Arts and industry groups such as the Australian Electrical and Electronic Manufacturers' Association and the Australian Information Industry Association. A draft should be delivered by mid-October and the final report should be ready for Cabinet approval by the end of the year. But DIST already has powers to make changes under the White Paper, and new CPE arrangements should be in place this month or next. A wider range of IT products will be covered by next March.

The BIE review has involved interviewing all 38 PDP/FTA participants and communicating with their overseas head offices with the aim of establishing how the companies allocate their global activities to different countries, and how the PDP/FTA schemes have impacted on their view of Australia as a place to invest in manufacturing and R&D.

Annual reviews and audits by DIST show that in calendar 1993, PDP/FTA activities generated exports of $1.2 billion and R&D activity worth $300 million.

Bernard Levy

One company bought a 155 Mbps cable system. The other one thought it did.

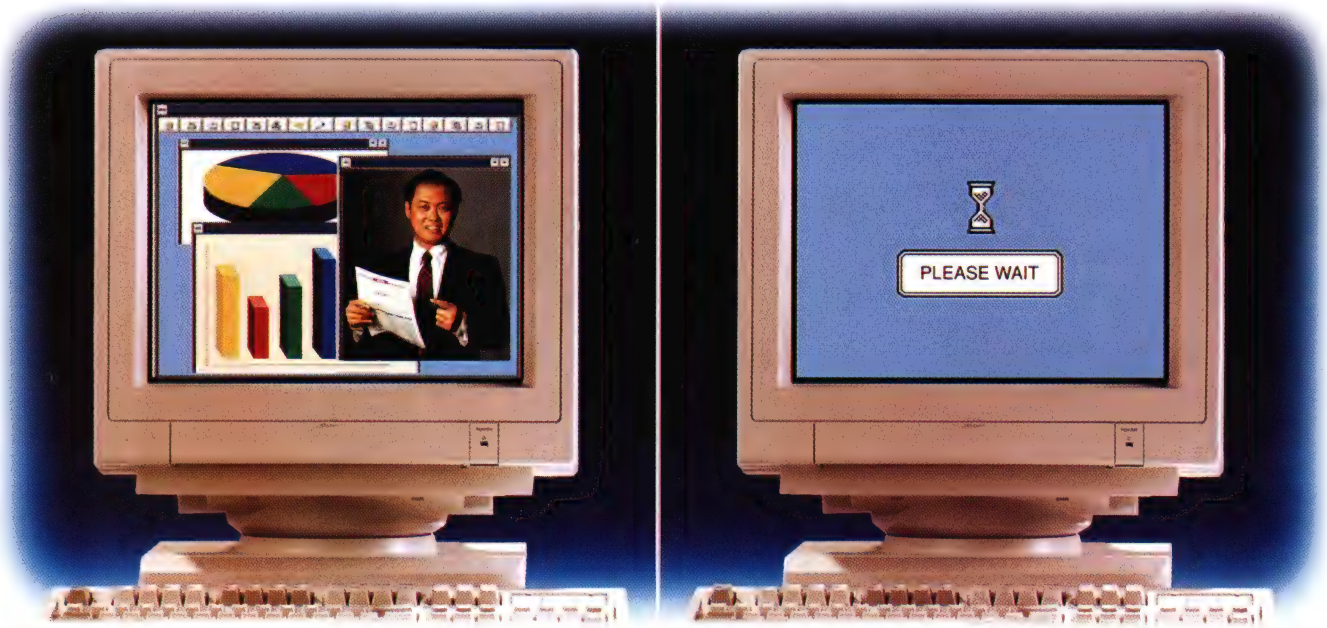

Many of today's cable systems promise Category 5 performance they can't deliver. Simply because they're not engineered to meet Category 5 demands.

That's not the case with AT&T SYSTIMAX® Structured Cabling Systems (SCS). They're engineered by AT&T Bell Laboratories to meet the strictest quality control standards. And *every* SYSTIMAX SCS product must meet or exceed EIA/TIA-568 Category 5 standards.

Part of our secret is in the installation.

That's why we handpick our SYSTIMAX SCS distributors, value-added resellers (VARs) and System Integrators (SIs). And why we train our VARs and SIs for weeks to meet our installation standards, from the maximum bend of the fibre to how much insulation can be stripped from copper cable.

All of this allows SYSTIMAX SCS to *deliver* speeds up to 155 Mbps. Support CD-ROM, Full-Motion Video, FDDI, SONET and ATM. And support even more advanced applications as new technologies emerge. Which means your network is virtually future-proof.

Your SYSTIMAX SCS certified installation also includes AT&T's 15-year Extended Product Warranty and Application Assurance.

So for peace of mind, insist on SYSTIMAX Structured Cabling Systems from AT&T, a world leader in data communications. For more information call us on 008 804 738.

AT&T SYSTIMAX®
Structured Cabling Systems.
The last cable system you'll ever need.

AT&T
Network Systems

For more information fax us on (02) 352 9279 or phone 008 804 738 or send this coupon to AT&T Australia Limited, Locked Bag No. 19, North Ryde, NSW 2113.

Name .....

Company .....

Address .....

..... Postcode..... Phone.....

AC

There's a *stack* of good reasons to *choose* HP Hubs.

Industry compatible

Affordable

yes NetWare Tested and Approved

HP, inventor of the stackable hub, now brings you the next generation. As you can see, there are now even more reasons to make sure you specify the market leader – HP.

The innovative, cost saving, new HP AdvanceStack hub is also standards based, and carries a lifetime warranty with a 5 year HP on-site warranty.

SPECIFICATIONS

- *Available in 12, 24 or 48 ports*
- *Manage up to 784 nodes with a single managed hub up to 185 metres apart*
- *Eavesdrop protection and access prevention security can be added with a management card*
- *Supports a range of media types including UTP, co-ax and fibre-optic with internal tranceivers*
- *Add bridge/routing and high speed technologies via the expansion slot*

For more information, see your nearest HP dealer or phone 13 1347. Or, to receive a fact sheet, phone our HP FIRST fax service on (03) 272 2627 and request ID number 7395.

HEWLETT® PACKARD

A Better Way.

JWT HPC 387

Microsoft's Daytona Dawn

Code-named Daytona, Windows NT 3.5 contains a swag of networking features. Enough for Microsoft to finally crack the network market?

Most people who have had to support Microsoft environments and applications have something of a love-hate relationship with the company and its products — what support specialist hasn't wanted to rewrite DOS with an axe or commit mayhem upon the people responsible for those truly informative unrecoverable application error (UAE) and other messages in Windows? Network support personnel, administrators and designers are no exception.

Microsoft has always built network operating systems which interact properly with their host operating systems. Unfortunately, they have never been complete; the Microsoft redirector, which was the heart of MS-Net, was almost bullet-proof (ill-behaved applications and TSRs aside) but it was slow, was only shipped with a clumsy, non-routable protocol stack (fortunately only to OEMs who knew better), didn't support peer-to-peer connections, was broadcast intensive and had no server operating system to support file and print sharing.

LAN Manager, the successor to MS-Net, cured the lack of a decent server and was much faster, but only supported peer-to-peer connections under OS/2, and was (until 2.1) tedious at best to tune and, although later versions introduced a TCP/IP protocol stack with a limited SNMP agent, was a glaring example of Microsoft's all but total ignorance of the realities of internetwork design and management. LAN Manager did not have any form of bridging or routing between IP networks, used a crude look-up table to resolve NetBIOS names and IP addresses (the LMHosts file) which could only be automatically updated in an internetwork environment via 'creative' use of LAN Manager's file replicator service. OS/2's basics like the user interface's look and feel, consistency of command sets and utilities like back-up also left something to be desired.

All of this — added to users' preference for the Windows 3.x interface, the need to build a truly portable operating system, and of course Mr Gates' vision of 'information-at-your-fingertips' — led Microsoft to create Windows NT, which has now been shipping for about twelve months.

Getting the Basics Right

Microsoft's aim with the first shipment of NT, version 3.1, and its advanced server version, NT-AS 3.1, was to get the basics right; a bullet proof OS engine, file system, comprehensive hardware support and complete set of utilities like backup, user and server admin, event log, etc. and above all an extensive, consistent and comprehensive API set to attract application developers.

To be credible in the marketplace, provide interoperation with the hugely successful Windows and Windows for Workgroups, provide a leading edge platform for application developers, and to avoid leaving its LAN Manager users high and dry, Microsoft had to put network plumbing into NT 3.1 and provide a server version (NT-AS 3.1). However, these network capabilities were clearly secondary to the primary objective of building an operating system able to run on several different hardware platforms and attract developers, and therefore users, in droves.

NT now ships for Intel, MIPS and Alpha-based machines and has been demonstrated on a PowerPC platform. Support for multiple processor platforms has been sufficient to get the attention of applications developers and vendors; however it is the clean, consistent and comprehensive API set NT presents which has resulted in those developers and vendors migrating so many applications to NT so quickly.

In particular, NT has received heavy support from vendors who produce statistical, database, engineering, graphic arts, etc. packages; applications which would normally be associated with Unix workstations. These vendors are used to supporting multiple processor architectures and, now that they have a comparable platform to work on, are keen to lift sales volume to the levels enjoyed by vendors of more traditional PC products.

Welcome to the Network

Having successfully delivered the basics in NT 3.1, Microsoft has now started to focus on NT's other aspects. While NT 3.5 — code named Daytona — does bring some new features to the basic operating system (see 'New Features in NT 3.5' on page 25) and several notable refinements, it is the in the area of networking that Microsoft has done the most significant work. Both Windows NT-Workstation (NTW — the new title for basic NT) and NT-Server (NTS — the new title for NT Advanced Server) have been rebuilt, from a network point of view, to make the deployment, maintenance and support of a large network based on NT as simple and cost effective as possible.

This is not to say that NT 3.5 is 'complete' in its range of network features — for example the IP router lacks support for wide area links via PPP (Point-to-Point Protocol) and SLIP (Serial Line IP) although RAS (Remote Access Services) can use either — however all the basic building blocks are now there and, when Microsoft's other enabling products such as SNA Server, Hermes, etc. are taken into account, an NT-based enterprise network, (rather than just a standard LAN interoperating with another more traditional environment), moves from the realm of future vision to that of practical option.

The most significant of the network oriented additions delivered in NT 3.5 is Microsoft's new 32-bit TCP/IP stack —

Figure 1: Windows NT 3.5's NetWare Interoperability

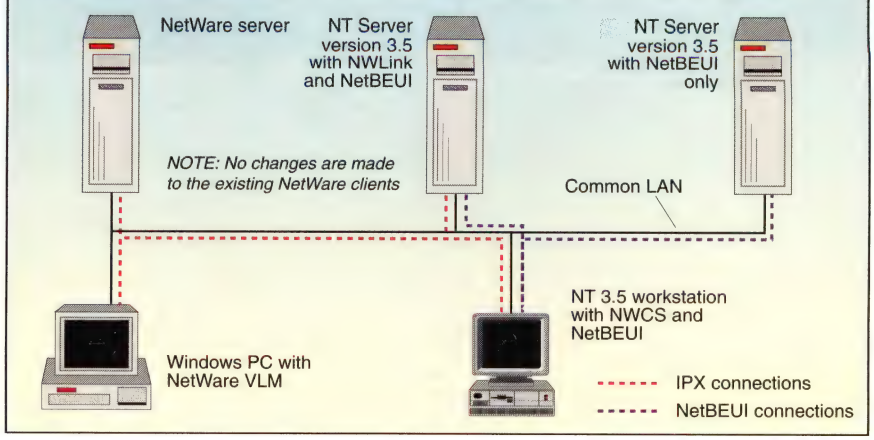

Market Watch

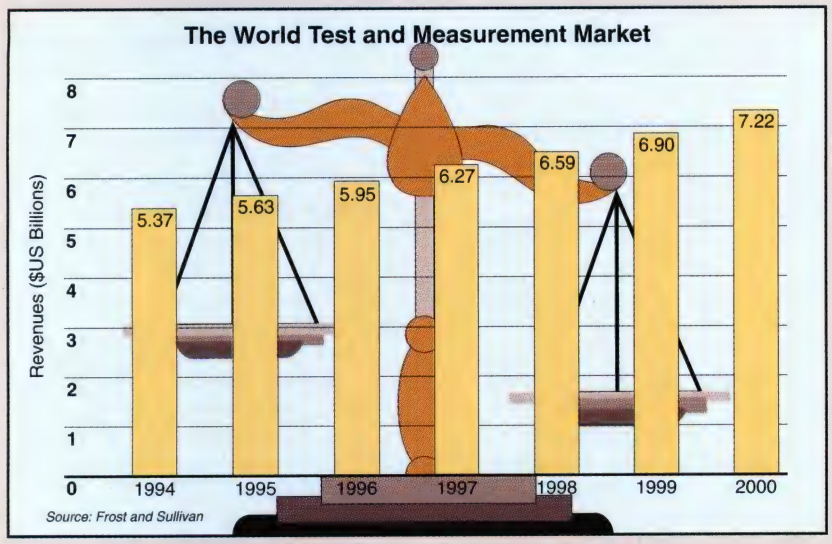

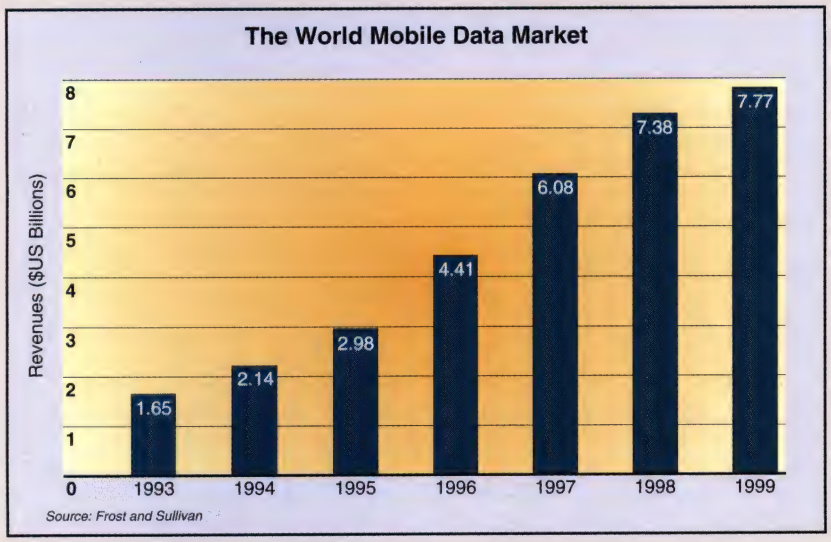

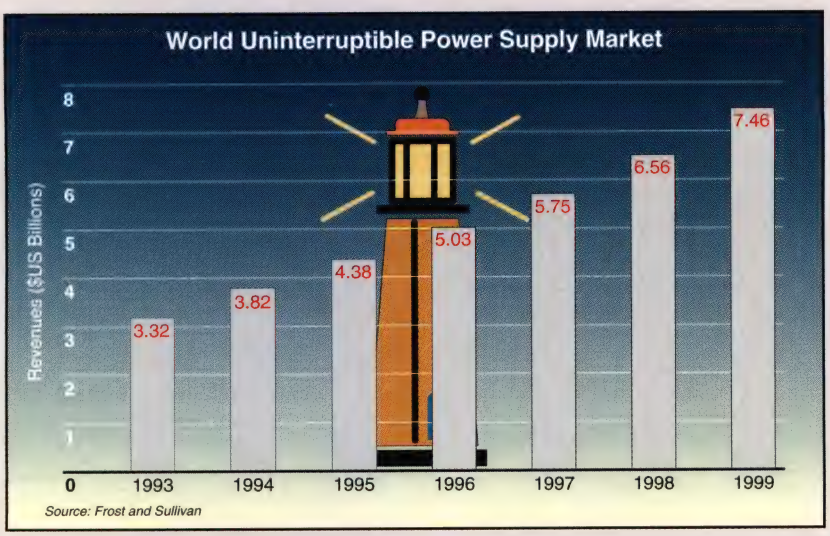

codenamed Wolverine — which has recently been delivered for the Windows for Workgroups platform. It is this new stack which makes the other IP related network enhancements in 3.5 possible.

Wolverine is in fact 'common code' which will run on Windows for Workgroups 3.11, NT 3.x and Chicago. Aside from being smaller, faster, containing more useful utilities (including an FTP Server in NTS), etc. than its predecessor, it supports the new DHCP (Dynamic Host Configuration Protocol) Service which, together with the WINS (Windows Internet Name Service) virtually eliminates the overhead of client address administration which has traditionally been associated with the deployment of TCP/IP in many PC environments.

Microsoft's implementation of DHCP conforms to the relevant RFCs (numbers 1533, 1534, 1541 and 1542) and therefore should interoperate with other non-Microsoft clients and servers. When a client (PC or otherwise) first boots it broadcasts a DHCPDISCOVER message, containing its MAC layer address and its network name on the network. A DHCP server will receive this message (even across a BOOTP relay router) and will respond with a DHCPOFFER message containing its MAC and IP addresses, an available IP address and subnet mask, along with an expiry date, for the client which issued the DHCPDISCOVER message. A client may receive more than one DHCPOFFER and having received one or more of them it sends a DHCPREQUEST message to the server whose offered configuration it chooses. This server then responds with a DHCPACK message assigning the IP address. When the client receives the DHCPACK message it becomes a *bound* DHCP client and can operate with that address until the expiry date of the address, at which time the request process is repeated.

DHCPREQUEST and DHCPACK messages also contain information about the IP addresses of a default gateway and both a WINS (Windows Internet Name Service) server and a DNS (domain name service) server if they exist. The DNS server would be a third party host such as a Unix box while the WINS server would be an NT Server, possibly the same one hosting the DHCP server.

WINS replaces LMHosts file, although this is retained for backward compatibility, and maintains a database of IP-to-Text machine names. For bound DHCP clients this mapping is automatically done during the binding process. During this process a NAMEREGISTRATIONREQUEST is generated to which the WINS will issue a positive or negative response. The former includes a time-to-live parameter, which has a similar function to the DHCP expiry date, and the later indicates that the name requested is not unique on the network, thus avoiding duplications automatically.

New Features in NT 3.5

NT Version 3.5 in both Server and Workstation versions will deliver the following new features:

- **Improved Set-Up Program**

Installation has been simplified. Upgrades from version 3.1 are virtually transparent, program groups, user and security definitions and network connections are automatically retained.

- **Reduced Memory Requirements**

Working set has been reduced to between 4 and 8MB depending on configurations. (NOTE: This is the amount of memory dedicated to the operating system *not* the amount required to run NT!)

- **Improved Performance**

Microsoft claim improved performance for both 32- and 16-bit applications; between 50 and 150% is claimed. Application load time has been decreased by over 50% over NT 3.1 say Microsoft. The author has noted significant increases in application speed, response and a perceptible decrease in load times while working with the beta software (which has several megabytes of debug code slowing it down).

- **Multiple VDM Support**

16-Bit applications run in their own memory space with full protection. This means that an UAE (unrecoverable application error) in one application will not shutdown all 16-bit applications, applications are more responsive and data protection is increased. OLE (object linking and embedding) and DDE (dynamic data exchange) are to be supported between VDMs (virtual DOS machines)

- **OpenGL Graphic Support**

Microsoft has licensed Silicon Graphics' Open GL graphics system. This means that high performance 3D graphics functions are included in NT 3.5's native API set. Developers of CAD/CAM programs will love this.

- **Integrated NetWare Connectivity**
- **Reverse engineered full NetWare Client**
- **New TCP/IP transport**
- **Common 32-bit TCP/IP stack with static router**
- **Dump Facility with Auto-Restart**

In the event of a fatal error memory registers are written to disk and the system can be automatically rebooted.

- **Long Filenames for FAT system**

File and directory names on the File Allocation Table (FAT) file system are no longer restricted to the 8.3 naming convention. NTFS naming conventions may now be applied to FAT volumes; names can be up to 255 characters and case is preserved. When long names must be displayed in 8.3 format they are concatenated for display only.

- **Account Lockout**

After a set number of bad password attempts (configurable by an administrator) an account is blocked. Even valid passwords for the blocked account will not be accepted until the account is manually enabled by an administrator.

- **Enhanced Hardware Compatibility**

NT has now been tested on over 2,300 Intel-based systems, 37 MIPS-based systems, 4 Alpha AXP systems and 15 SMP (Symmetric Multi-Processor) systems.

- **Enhanced Multimedia Support**

Video for Windows and 32-bit drivers for video compression using CinePak are included.

- **Plotter/Scanner Support**

Drivers for HP plotters and various scanners are now included in the base NT.

- **Multiple Network Support Under IPX**

Includes SPX II support

- **NetWare Compatible Gateway**

Allows Users to access NetWare servers without running dual stacks. Intended as a transition tool.

- **DHCP Service**

Dynamic Host Configuration Program assigns IP addresses dynamically from a central pool.

- **Windows Internet Name Service**

WINS maps text names to IP addresses automatically.

- **Enhanced RAS**

The Remote Access Service now supports 256 concurrent users per server. SLIP, PPP and Novell NetWare over PPP connection are also supported.

- **Win-16 Administration Tools**

NT servers can be administered from Windows 3.x or Windows for Workgroups machines.

- **Remoteboot (RPL)**

LAN Manager (version 2.2) equivalent support for diskless workstations — i.e. DOS and Windows 3.1 clients.

- **TCP/IP Printing**

Allows NT 3.5 servers to print directly to printers connected to Unix hosts or TCP/IP printers directly connected to the network. Also acts as a gateway for servers which do not have TCP/IP installed allowing them to use Unix-connected and TCP/IP printers.

- **Dynamic File Compression**

Provides the capability to transparently compress files on a per-file or directory basis within an NTFS volume. Microsoft claim better performance than DoubleSpace and that there is provision for extension by third parties.

NOTE: This list has been prepared on the basis of the Beta release of Windows NT 3.5 and may change prior to final release. Microsoft intends to have both NTS and NTW in stock in Australia at the end of September.

Licensing — Both NTS and NTW will be licensed on a per CPU basis. NTS will be functionally unlimited in the number of users supported, as is the present NT Advanced Server 3.1, however user licences will be legally required and will be available on a per client basis. For example, a site with 200 users and 4 servers will buy four copies of NTS (or 1 copy and 3 licences) and 200 user licences. Upgrade packages will come with 250 user licences. Microsoft is moving to this licensing model for all its system products.

WINS information is updated across the entire internetwork on a scheduled basis which, after initial set-up (which has to be done with some care), requires no intervention from administrators. Moreover, because IP addresses are not irrevocably bound to MAC layer addresses, DHCP clients can move to physically different locations without manual reconfiguration. The Sydney-based user can plug his/her laptop into a network port in the Melbourne office and, given an appropriate software set-up, work as normal via a router connection to the Sydney servers (or a LAN link to the Melbourne ones) and, since DHCP will not reallocate inactive addresses until it has run out of unused ones, it is very likely that the user will be reassigned their original IP address when they return to their Sydney office.

The downside is that security schemes based on IP addresses will be negated by DHCP since a mobile host will be assigned an address on a valid local subnet whenever they become active on the network. The operation of MAC layer security schemes however will be unaffected by DHCP. In sites where such schemes are employed mobile users will have to request a network port be assigned for their use — which is not a bad idea.

From the user's point of view they will — at last — be able to log on to a network at any physical location and still see all the active resources on the network via the familiar Windows 'browser' dialogue box. All this happens with a drastic reduction in broadcast activity on the network.

A Static Approach to Routing

Daytona's server version also introduces a static IP router designed to route IP traffic between networks bound to NDIS 3.0 devices — i.e. LAN adaptors or third party devices which provide wide area connectivity via an NDIS 3.0 interface. Static routers (which direct traffic along routes manually set and do not change dynamically in response to altered network link states) have largely been replaced by dynamic routers which employ protocols like RIP and OSPF to communicate route information which enables a mesh topography internetwork to be fault tolerant (in so much as faulty links are concerned) and to load balance across multiple links.

The downside to the intelligence of a dynamic router — which is minor in most cases — is that they are complex devices which have a relatively high processor overhead and therefore are generally slower than a static counterpart, can be broadcast intensive — especially when RIP is used — and can be quite complex to configure. If two networks are connected via a single link or several routers are interconnected via a mesh topography which never — or at least very rarely — changes then there is some case for the use of simple static routers.

This was an issue which generated hot debate in the late '80s when dynamic routers were first being introduced. One of the points then in favour of the static router was that there were some quite large internetworks built around static routers in operation at the time. The point is that you can do quite a bit with simple static routers — especially in a local or campus environment — though nobody today would argue that a dynamic router does not, in general, do a better job than a static one.

Figure 2: Windows NT 3.5's NetWare Gateway

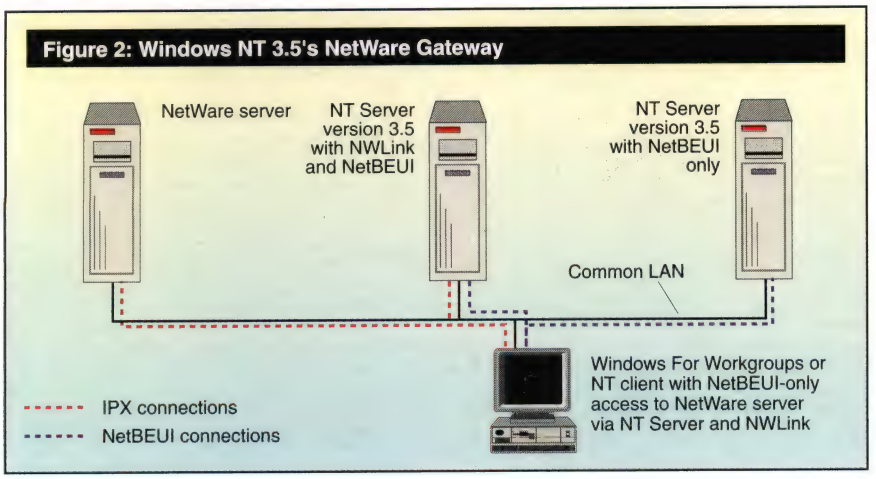

What Microsoft have provided, in recognition that some form of router is an essential component of any practical network server, is a simple static router with a simple — in IP terms — user interface. This allows administrators of small to medium sized network sites or networks comprising small remote sites to install fire walls and control traffic flow by segmenting the system into a number of logical LANs with well defined linkages. You have to think about addressing schemes a little more carefully than you might if you buy a purpose built router, but you won't have to pay for the router.

A number of companies are working on more sophisticated routing products for NT although, at the time of writing, only Eicon has delivered. The Eicon IP Router for Windows NT works on version 3.1 for the moment, supports all Eicon's cards, supports PPP, RIP, is compatible with purpose built routers and all Eicon's other IP routers. Eicon has also demonstrated a product called WAN Services for Windows NT which is designed to provide X.25 connectivity for Microsoft's Enterprise Messaging Server (which will be released as the Microsoft Exchange Server).

The NetWare Connection

Microsoft has also upgraded NT's NetWare connectivity by shipping a reverse engineered NetWare client (which supports SPX II) with both NTW and NTS, and a NetWare gateway called NWLink with NTS. The former allows an NT station to access a standard NetWare server via the standard Windows interfaces just as a DOS/Windows or Windows for Workgroups-based NetWare client can (see Figure 1 on page 23). This means that those who wish to retain a NetWare-based environment can add NT-based applications servers (Notes, SQL Server, SNA Server, etc.) or workstations to their existing networks with little or no impact on their NetWare environment.

The NetWare gateway is intended as a migration tool and allows NT Server clients which are not running a NetWare stack to access files stored on and printers attached to a NetWare server without running an IPX/SPX stack (see Figure 2). Microsoft, not surprisingly, consider NT to be a better platform than NetWare but, at least in Australia, it seems to see competition with Novell (and others) as almost incidental to having the market accept NT as a real alternative to minis and the like as a serious platform for core business applications.

Nevertheless, NT in combination with products like Windows for Workgroups 3.11, Office 4.3, SQL Server and Visual Basic can provide an extremely cost effective platform for custom systems for small sites — 30 users or less — as well as large sites and it is these small sites which are the bulk of Novell's market in this country.

For those using SQL Server there are some compatibility issues which should be considered when moving to NT 3.5. Microsoft has ended its development agreement with Sybase but both parties have agreed to maintain compatibility between their implementations of the Transact SQL language — to the version 4.2 level. The result is that the new features in Sybase's System 10 product and Microsoft's current version 4.21 SQL Server are incompatible. These incompatibilities are, at this stage, minor but they are likely to grow.

Microsoft SQL Server 4.21 will not run on NT 3.5; it requires a patch — version 4.21a. A new version of SQL Server is due for release some months after the release of NT 3.5 (due to the need to test it against the final release of NT). It's code-named 'SQL 95' and along with its management interface — code-name 'Starfighter' — will require NT 3.5. It will not run on version 3.1, so plan ahead. Note also than Microsoft will only support SQL Server on the OS/2 platform for a further two years.

In addition to TCP/IP and IPX/SPX support, the base NT Server product also provides a DLC (data link control) protocol stack for printers, 3270 emulators, etc., and Services for Macintosh. NT Server Services for Macintosh, once installed, permit Mac workstations to transparently access NT servers as full clients, sharing both file and printer resources; however security is limited to that given by the standard Apple system. If administrators wish to extend NT's security regime to Mac clients then an optional Authentication Module must be installed. NT Server Services for Macintosh includes a full AppleTalk Phase II router.

Remote Considerations

Microsoft has also made significant improvements to NT's Remote Access Service (RAS). RAS now supports PPP, SLIP (effectively acting as an IP router) and IPX over PPP (effectively acting as an IPX router) in addition to NetBEUI (in which case it acts as a NetBIOS bridge). The great advantage of this approach is that, by using PPP or SLIP rather than the previous proprietary scheme, NT RAS clients can use any PPP or SLIP access server to access a network. This provides the potential for direct, dial-up access to, for example, the Internet via a simple terminal server which supports SLIP or a router which supports PPP and modem connections.

The RAS server has also been upgraded to support 256 concurrent sessions rather than 64 as in version 3.1. Anymore than a handful of concurrent dial-in users would almost certainly require a dedicated RAS server and alternatives such as a PPP or SLIP server or the like of 3Com's AccessBuilder (see 'Making the Office Accessible' on page 50) might be a more cost effective solution.

So Far, So Solid

This article has been written using Microsoft Word for Windows 6.0a on a 486 running NT 3.5 attached to an Ethernet network with NT 3.1 servers and Windows for Workgroups PCs — all 486s. Daytona has been tested over several weeks with all the applications in MS-Office 4.3 Professional and a few other odds and sods. 'Tested' ranges from file opening, saving and closing to turning the power off in the middle of moving several megabytes of files from one directory to another.

Through all this Daytona has proved to be solid and reliable; it interoperates seamlessly with the other variations of Windows on this LAN and the installation from CD is positively boring — I'd have fallen asleep if there had been time. Assuming Microsoft delivers NT 3.5 in late September as planned, and components such as SQL Server soon after, I suspect that it will be a hard product for administrators of small and large networks alike to ignore. While there is still plenty of scope for improvement — e.g. direct WAN support in the IP router, say via com ports and modem, or a fax gateway as in Windows for Workgroups — Microsoft might just have gotten it right this time.

Graeme Le Roux is a Director of Moresdawn Pty Ltd (Bundanoon, NSW) and specialises in local area network consulting services.

THE NATURAL CHOICE IN NETWORKING

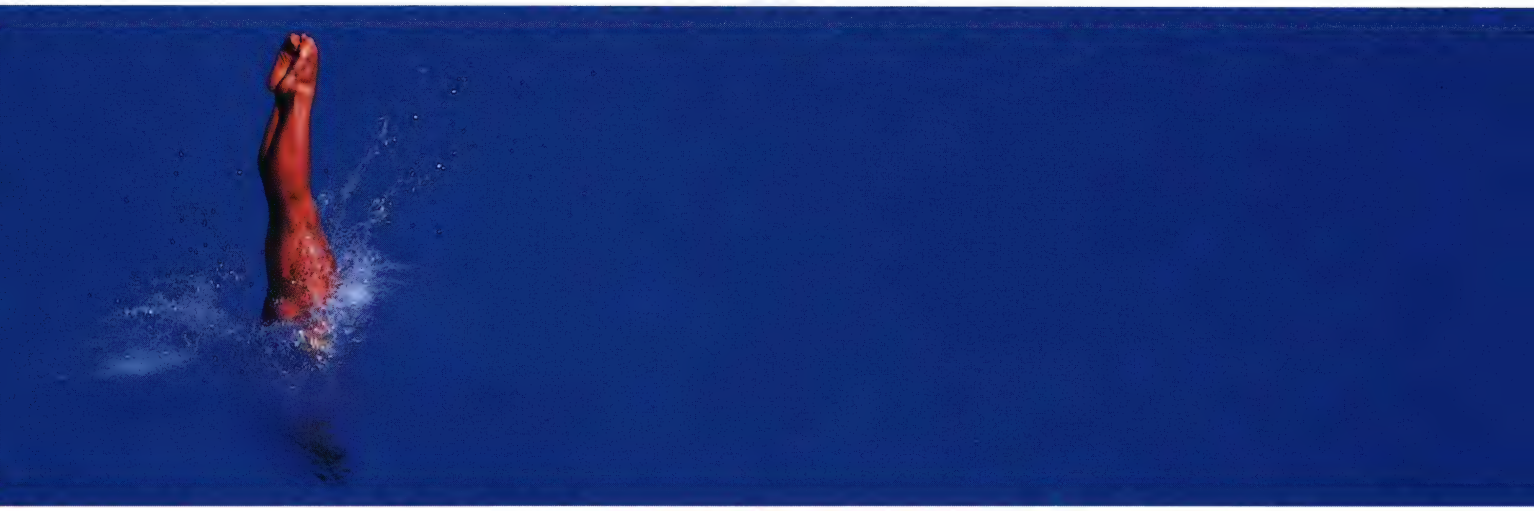

You'll make a smoother entry into ISDN with Racal Australia.

Who wouldn't like the benefits of ISDN? You know, instant call set-up, high speed, lower costs.

But perhaps you're uncertain about putting a toe in the water.

Well, now you can safely take the plunge with Racal.

We have the internationally approved and supported products - like terminal adapters, multiplexers and back-up units - with the expertise to make ISDN easy.

And we can offer full network management to simplify running an ISDN-based operation.

But there's another way we can make ISDN even more attractive.

We'll help you look at all possible methods of reducing your initial investment in lines and equipment.

To find out more, call Racal Australia on (02) 936 7000, or post or fax the coupon.

It really is time to dive in.

Racal Australia Pty Ltd, 3 Powells Road, Brookvale, NSW 2100.
Telephone: (02) 936 7000. Fax: (02) 936 7036.

☐ **Tell me more about Racal Australia's family of ISDN products.**

Please send me information on Racal Australia's other products:

☐ **WANs** ☐ **LANs** ☐ **Security** ☐ **Recorders** ☐ **Instruments**

Title Initials Surname

Position

Company

Address

Post Code

Telephone Fax

Communicating through technology RACAL

Channel Extension

Changing Channels for Corporate Connectivity

An entirely new set of channel technologies has emerged that will no less than transform the role of the channel in the enterprise. The muscular new kids on the block are Escon, which replaces bulky copper with fibre, and Fibre Channel, which provides an open systems standard for a variety of devices.

Channels, the fastest way to get data into or out of a processor, are typically thought of in terms of mainframes. These critical data pipes usually come in three general configurations:

- *Host-to-host* communications (often referred to as channel-to-channel) encompasses such functions as file transfers between processors. Since these tend to be bulk-data-oriented, they involve very large block sizes and require substantial bandwidth;
- *Host-to-peripheral* communications covers the control that hosts exert over peripherals, as well as data being transferred between the two. Since peripherals vary considerably, the type of traffic associated with different devices varies as well. For example, storage-oriented peripherals like tape drives generate traffic patterns similar to the channel-to-channel exchanges in a file-transfer operation, while communications-oriented devices like terminal controllers have smaller block sizes and are more sensitive to response times;
- *Channel-to-LAN* communications is an emerging set of applications in which standard protocols like TCP/IP form a common language for various platforms while channel gateways establish physical connectivity between the two. In many respects, this opens a new role for the mainframe on internetworks and client-server applications.

In general, these sorts of channels tend to be very fast, capable of handling much more data than their LAN counterparts of the same generation. They also are designed to be highly resilient, since a channel failure can bring a mission-critical application to an abrupt stop.

IBM Country

The most common implementation of channels is in the IBM mainframe marketplace. IBM 360 and 370 series hosts and all of the plug-compatible mainframes use a channel specification defined by the US Federal Information Processing Standard (FIPS-60) commonly known as the bus-and-tag channel interface (see the table below). This spec calls for two copper cables: the bus handles the data, and the tag is used for control information for special sequences (such as the selection sequence, which establishes the initial connection to a device).

Among bus-and-tag's strengths are wide implementation and a relatively high data rate: 1.5Mbyte/s to 4.5Mbyte/s. But on the downside, it is very bulky, and is limited to 400-foot runs (about 121 metres).

A parallel on the workstation side of the market, and developed at roughly the same time as the bus-and-tag channel, is the small computer systems interface (SCSI). Like bus-and-tag, SCSI connects processors (in this case, primarily workstations) to other processors or peripherals. SCSI also enjoys broad acceptance in the market. It suffers from some of the same distance limitations as bus-and-tag.

Although analogous to the master-slave relationship in mainframe bus-and-tag communications, SCSI implementations use a different terminology. In SCSI, the device that controls the operation is referred to as the initiator, usually a host. The device that acts on the communication is referred to as the target, usually a peripheral.

Although SCSI is typically thought of as a single, open specification, it actually has a number of variations that affect everything from the connector and cabling to the electrical properties of the signal flowing over the cable. One dimension of these variations is found in the signalling characteristics that flow over the SCSI cable, which in turn determines its range. Single-ended SCSI has a maximum distance of six metres; differential SCSI can run 25 metres. Single-ended is cheaper to implement, so it is usually found within a cabinet or for system-to-peripheral communication. Differential SCSI is the preferred choice for system-to-system communications.

Data rates also vary across SCSI implementations. The original spec called for a maximum data rate of 5Mbyte/s, which is now called slow SCSI. A subsequent upgrade, called fast SCSI, provided for synchronous data rates in excess of 5Mbyte/s.

Another SCSI variation relates to the number of pins on the connector and how many of them are used in synchronous data transfers. The original SCSI spec called for an 8-bit transfer, now called narrow SCSI. But along came wide SCSI, which defined formats for 16 or 32-bit data transfers.

These characteristics are typically combined to form common SCSI implementations. The original SCSI, called SCSI-1, combined slow and narrow for an effective data rate ranging from 2.1Mbyte/s to 5-Mbyte/s. SCSI-2 provides for fast and/or wide data transfers, which again substantially improved data throughput. SCSI-2 narrow implementations range from 4.2 to 10-Mbyte/s, and SCSI-2 fast/wide implementations run up to a theoretical maximum of 40Mbyte/s.

The New Generation

Now, two technologies that overcome the inherent limitations of bus-and-tag and SCSI have made their appearance.

On the mainframe side, IBM announced in 1992 a new channel architecture for its

Choosing a Channel

	BUS AND TAG	SCSI	ESCON	FIBRE CHANNEL	HPPI	FDDI	ATM
Topology	Multidrop	Multidrop	Circuit switched	Circuit or packet switched; ring or point to point	Circuit switched	Token passing ring	Cell switched
Media	Copper	Copper	Fibre	Independent	Copper	Fibre	Independent
Data rate	1.5Mbyte/s to 4.5Mbyte/s	2.1Mbytes/s to 40Mbyte/s	17Mbyte/s	133Mbps to 1Gbps	800Mbps	100Mbps	44.7 to 155Mbps*
Maximum distance	121 metres	25 metres	20km	10km	25 metres	2km between stations; 200km maximum ring diameter	Unlimited
Packet/frame/ cell size	N/A	N/A	1,035-byte frame	2,148K frame	4.3GB packet	4,500-byte packet	53-byte cell

*Defined speeds range up to OC-4 (approximately 2.5Gbps)
ATM = Asynchronous Transfer Mode
Escon = Enterprise System Connection
N/A = Not Applicable
SCSI = Small Computer Systems Interface
HPPI = High-Performance Parallel Interface

The ATM Truth.

Newbridge Delivers End-to-End, Managed ATM.

ATM is the only technology that can work seamlessly across an enterprise – from the desktop to the wide area network. With our VIVID products in the LAN and ATMnet products in the wide area, you can capitalize on the seamlessness of ATM to achieve an end-to-end solution that meets all your bandwidth requirements. Even better, you can manage your network with a single system – the first network management system to integrate ATM, Frame Relay and conventional multiplexer management!

Start building your ATM network with Newbridge today and you can choose from a complete range of interfaces including Ethernet, LATM, T1, E1, T3, E3, OC3, STM-1 and video.

Whether you start with the VIVID line for LANs or the ATMnet family of WAN products, you can be sure of getting the performance and scalability you need with comprehensive network management. Intrigued by the truth? Call Newbridge.

Newbridge Networks Australia Tel: 02 255-7748 Fax: 02 247-2417

NEWBRIDGE

In the Americas, 1 800 343-3600 or 1 703 834-3600 | Europe, Middle East and Africa, 44(0) 633 413600 | Asia Pacific 1 613 591-6300.

Athens • Atlanta • Auckland • Bahrain • Beijing • Boston • Brussels • Buenos Aires • Caracas • Chicago •Copenhagen • Dallas • Denver • Frankfurt • Geneva
Hong Kong • Johannesburg • London • Los Angeles • Madrid • Mexico City • Miami • Minsk • Montreal • Moscow • New Delhi • New York • Paris
Rio de Janeiro • Rome • St. Petersburg • San Francisco • Santiago • Seattle • Shanghai • Singapore • Sydney • Tokyo • Toronto • Vancouver • Washington

MATRIX NB901

Figure 1: More From Less

Escon, IBM's fibre optic successor to copper bus-tag-channels, boasts a greater data rate: 17 megabytes/second vs. 4.5 megabytes/second. The technique is gaining acceptance among mainframe data centres.

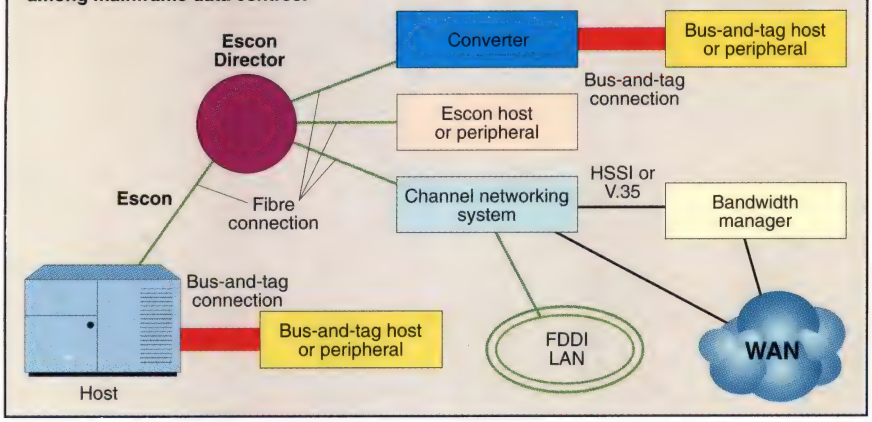

390 series called Escon (Enterprise Systems Connectivity). Rather than bus-and-tag's bulky copper, Escon channels use fibre optic cabling, making them considerably smaller and faster. Escon's data rate of 17Mbyte/s is substantially greater than the 4.5Mbyte/s maximum rate of bus-and-tag.

Another difference is that IBM, though keeping Escon proprietary, is actively supporting its development by third parties, mostly by selling Escon chip sets.

An important improvement offered by Escon is the Escon Director, a high-speed circuit switch that allows a single Escon channel to communicate with many different devices. Circuit switches like the Escon Director offer an important advantage over other networking technologies because each path through the switch has access to the full rated bandwidth of the circuit. The switch operates by setting up a logical connection between two ports and maintaining it all through the transfer. At the same time, the switch can provide the same kind of connection between other devices.

In contrast, shared media networks force nodes to compete for a fixed pool of available bandwidth. That means that throughput available for a given device drops as the number of active sessions increases. This gives circuit switches like the Escon Director an aggregate capacity far exceeding the bandwidth associated with shared media networking technologies. Escon Director tops out at about 6Gbps.

Range limitations of the bus-and-tag channel also are partially relieved by Escon, which can extend up to 20km without repeaters, 60km with repeaters. The effect of this improvement, however, only has value for certain applications. The Escon architecture itself does not define a means of connecting to wide-area circuits, meaning that the organisation must be able to lay cabling between the host and the target device. If the channel must run outside a building or campus, that involves right-of-way issues.

Word to the Wise

By moving from copper to fibre, Escon dramatically changes the physical properties of the host channel. But it uses exactly the same device command set, called channel command words (CCWs), as do bus-and-tag channels. In fact, this CCW structure has been in place for the most part since the 360-style mainframes were developed in the 1960s.

The physical properties of fibre enabled IBM to provide a number of important improvements over bus-and-tag, including increased speed, greater range, and lower interference. But by maintaining the CCW structure of earlier generations, virtually no significant changes had to be made to the host channel interface software or the peripherals offered by IBM. In large part, migrating to Escon means adding a physical Escon channel and connecting it to an Escon-compatible device. The drawback, here, is that Escon maintains the proprietary nature of IBM channels.

As another aid in the migration to Escon, IBM offers a device called the Escon converter that takes advantage of the consistent CCW structure. Because of this, the converter needs only to change the bus-and-tag's parallel/electrical channel signals to the serial/optical signals of Escon and vice versa: no additional protocol conversion is necessary. This relatively easy procedure allows customers to begin their migration by first investing in Escon hosts without forcing them to immediately upgrade their bus-and-tag peripherals.

Despite the advantages, sales of Escon to date have been sluggish. All 390 models being shipped by IBM offer both Escon and bus-and-tag channels, and more can be added. The main limiting factor is that the vast majority of peripherals in the field or on the market use bus-and-tag, reducing the usefulness of Escon until more Escon-equipped peripherals become available. IBM's 3390/3990 DASD (direct access storage device) units and 3490 tape systems are among the first on the market. More are sure to come in the near future.

Fibre Channel

Much as bus-and-tag and SCSI were being developed at roughly the same time, Escon also was developed in tandem with a channel architecture designed for the high-performance, open systems market: Fibre Channel. Defined by ANSI Technical Committee X3T11, Fibre Channel is intended to relieve network bottlenecks created by the

Quietly doing your job

Nortel's IBDN® cabling system keeps pace as your networking requirements change. IBDN® provides many benefits: performance, high reliability, flexibility, cost-effectiveness and the capacity to evolve.
No matter what size your company is, Nortel has a solution. For additional information call CableNet, a Nortel Master Distributor for IBDN® and FOTS products.

CableNet
6/30 LEIGHTON PLACE, ASQUITH 2077
TEL: (02) 482 7300 • FAX: (02) 482 7311

nt nortel

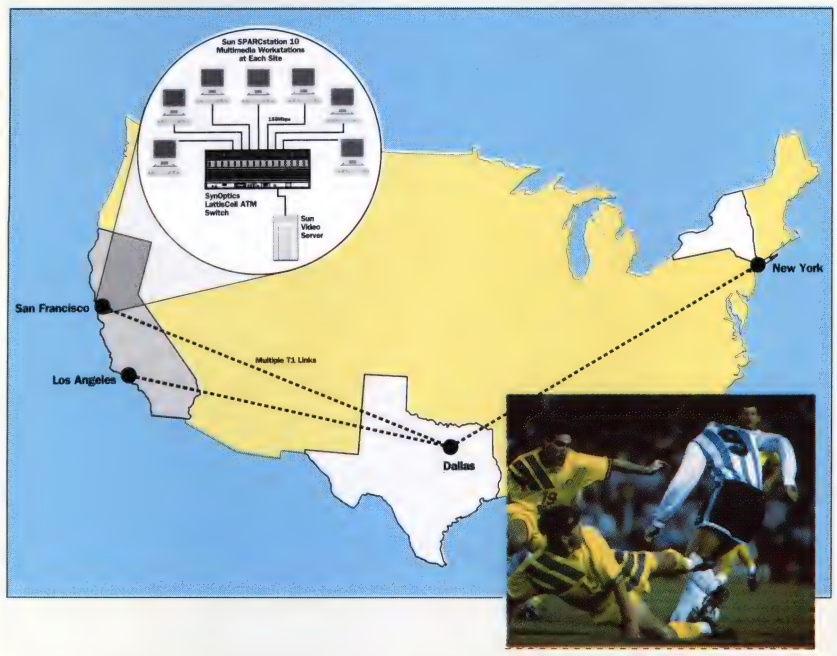

SynOptics Scores at the World Cup

SynOptics ATM products were used to link the nine US World Cup '94 sites to provide on-line services, such as real-time video and audio clips of match highlights and news and information, to soccer fans attending games. Part of the World Cup Multimedia Project, SynOptics LattisCell ATM switches were linked to 'multimedia kiosks' at World Cup venues in Dallas, Los Angeles, New York and San Francisco. Each multimedia kiosk was equipped with full-motion, full-sized NTSC (640 x 480) 24-bit colour video systems, built around Sun SparcStation 10 workstations linked by 155Mbps fibre optic connections to LattisCell switches and Sun video servers. Services available to soccer fans included full-motion colour video clips of matches at other venues, video clips of individual player highlights, interviews with coaches, and on-line information such as game and team statistics.

rapid increase in processing power of workstations. The standard provides ultra-fast 133Mbps to 1Gbps local connectivity.

A unique characteristic of Fibre Channel is its topology-independence. The technology can be implemented either as a point-to-point connection or to provide any-to-any connectivity. In fact, Fibre Channel defines two any-to-any network topologies: switched fabrics and arbitrated loops.

Switched Fibre Channel can provide high throughput and any-to-any connectivity among a large number of nodes. Like the Escon Director, Fibre Channel switches can switch circuits to deliver very high aggregate data rates to all devices on the network. Unlike Escon, Fibre Channel switches also can use packet switching to handle packet-oriented data, such as LAN transmissions.

In arbitrated loops, network nodes are connected in a ring topology similar to that used by FDDI. But Fibre Channel rings allocate bandwidth based on arbitration, meaning that the two nodes involved in a given communication get full access to the ring's bandwidth until they are finished, forcing any other communicating pairs to wait. The advantage to arbitration is that every communication has the opportunity to access the full rated bandwidth of the channel; the drawback is that as the number of nodes on the ring increases, so will the chances of collisions in which communicating pairs will have to wait for access to the network.

Another aspect of Fibre Channel flexibility can be seen in the variety of service classes defined as part of the specification. Class 1 connections represent a dedicated circuit between two devices. They offer several advantages: short connection setup time, long time available for data transmission, and guaranteed delivery of frames. Class 2 connections use frame switching to share the network but do not guarantee that the frames will be received in the order they are sent. Class 3 connections are analogous to typical network technology datagrams, in that destination nodes do not acknowledge receipt of a transmission, making the transmission vulnerable to network congestion.

Like other standards-defining bodies, the ANSI Technical Committee X3T11 uses a layered approach for defining the Fibre Channel specifications. The physical layer, defined by FC-O, provides for various media implementations including copper, multi-mode (LED transmission) fibre, or single-mode (laser transmission) fibre cabling. Data rates range from 133 to 1062-Mbps; distances range from a few kilometres up to 10km. Other layers include FC-1, which defines encoding/decoding control sequences; FC-2, which defines the framing protocol; and FC-3, which defines common services.

Action Layer

But the real action is taking place at FC-4, the layer that maps operations (framing and/or signalling structures) defined by other technologies into Fibre Channel framing and signalling layers.

FC-4 definitions that have already been designated for development include IP, IPI-3, SCSI, HPPI (high performance parallel interface), and the bus-and-tag (block multiplexer) channel. In addition, other FC-4 specifications are being debated, most notably ATM; direct memory access; various LAN technologies such as FDDI, Ethernet, and Token Ring; and native bus interfaces such as SPARC and Microchannel. This ability to simultaneously run multiple network and channel command sets over one media demonstrates Fibre Channel's protocol-independence, and highlights its usefulness for a variety of network and channel applications.

Taken together, all these concepts — high bandwidth, flexible network topology, support for a variety of high-level protocols, and defined interoperability profiles — make Fibre Channel well suited to a wide variety of environments.

At present, Fibre Channel technologies are primarily found in scientific and high-performance computing markets, but many vendors say they plan to develop an interface. In fact, IBM has developed one for its RS-6000 workstations, and other vendors, including Hewlett-Packard and Sun Microsystems, have products under development.

Cray and HPPI

High-performance computing platforms also have experienced an evolution of channel technologies. In this case, the heritage channel is the Cray Channel, a parallel channel found on most super computers developed by Cray Research Inc. through the YMP models. This channel supports a data rate of 12Mbyte/s running TCP/IP for FTP (file transfer protocol) applications and IPI-3 for communicating with peripherals.

Various high-performance computing applications now use a successor to the Cray channel: the high performance parallel interface, or HPPI. It was designed in the late 1980s, a time when many scientific and government initiatives had a pressing need for increased bandwidth but could not wait for fibre optics to become well-established. As a result, the HPPI specification relies heavily on proven, widely available tech-

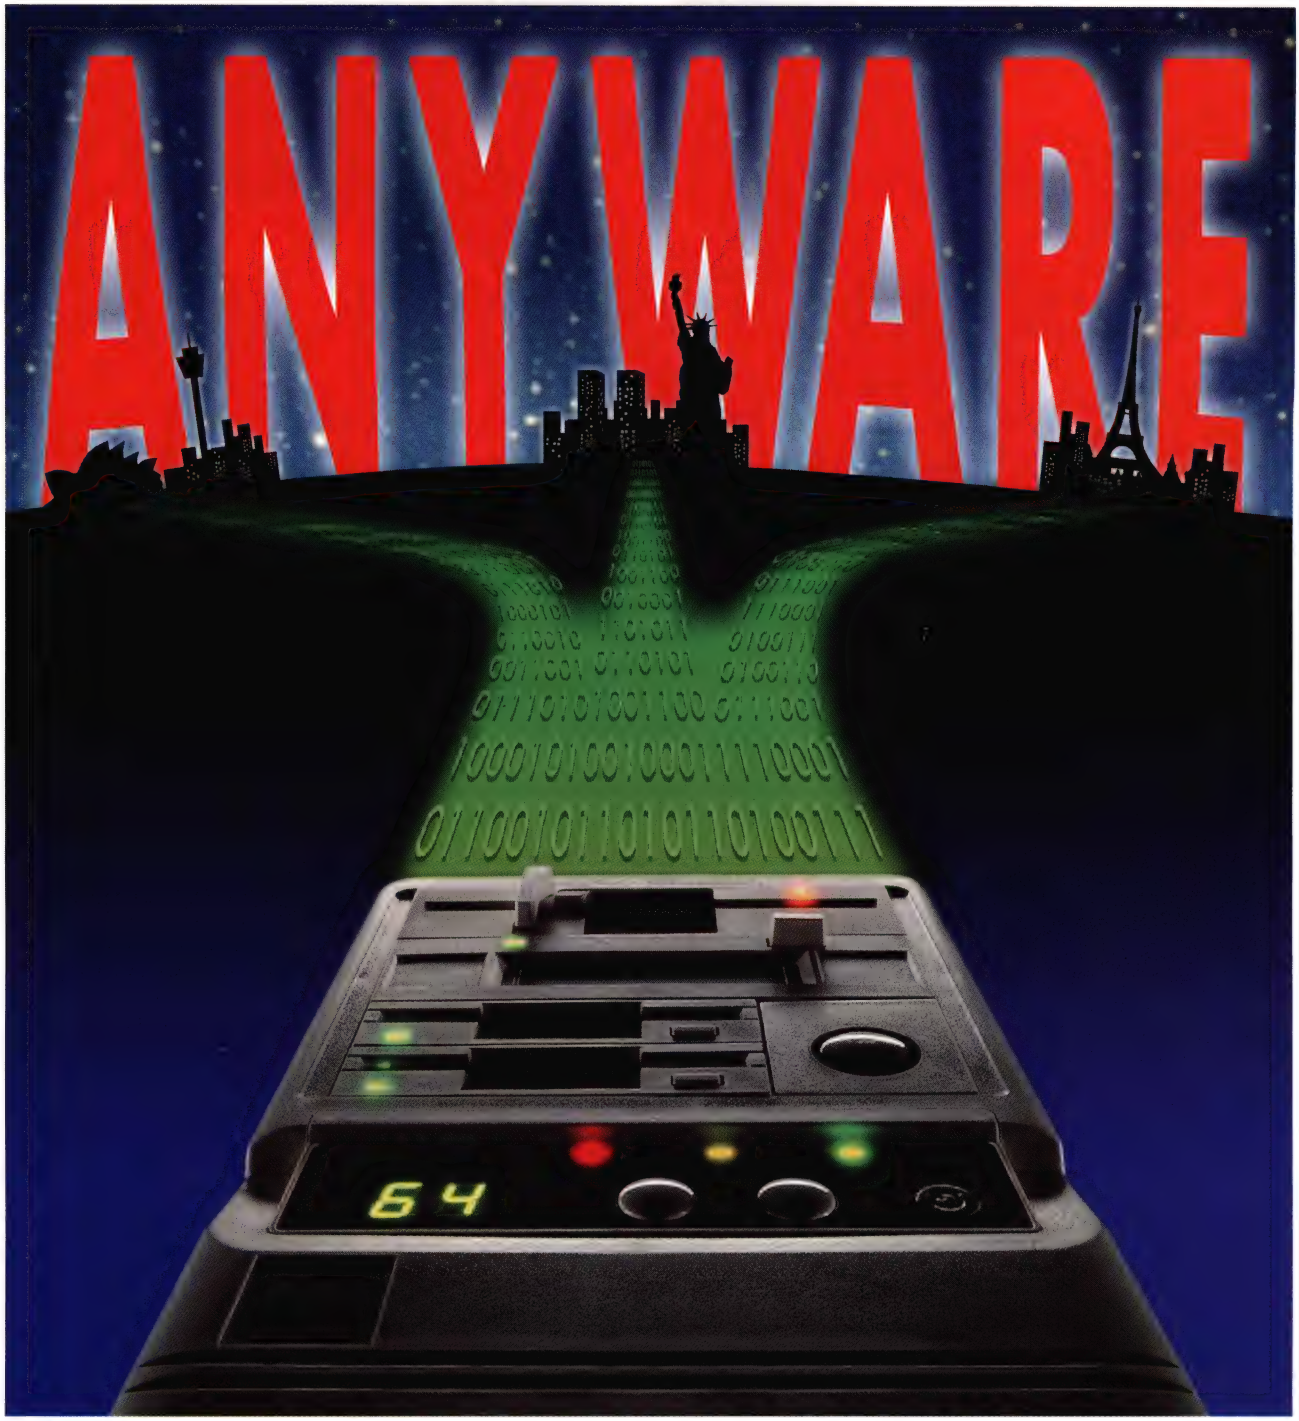

Connect NetWare® with ISIS AnyWare™ and save a fortune on the way!

Connecting your LANs together has never been easier, or so affordable. ISIS AnyWare™ Remote LAN Access products provide on-demand, high performance, transparent and secure connections between NetWare® servers and any other LANs. AnyWare, Anytime, Anyplace!

If you want to know more about our wide range of affordable, reliable, high performance and easy to use Remote LAN Access solutions, contact us for a copy of our *Guide to Remote LAN Access and ISIS*, or talk to one of our product specialists.

Contact our Asia Pacific headquarters in Melbourne on Tel. +61 (03) **629 3800** Fax. (03) 629 5791
or our Sydney sales support organisation, TeleKNX on Tel. +61 (02) **954 5611** Fax. (02) 954 5190

TeleKNX DISTANCE NETWORKING

ISDN Powered by ISIS

yes It runs with NetWare

NETWORK DESIGNERS KNX
LEADERS IN REMOTE LAN ACCESS

Concept, design & imaging: Multicomm Productions +61 (03) 500 9907

The next time your network needs a more advanced hub, you'll already have it.

For a truly intelligent way to network, why not choose a hub advanced enough to be its own replacement? A hub designed to easily accommodate any change your growing network requires. A hub flexible enough to allow you to customize and configure your network as it grows.

It's 3Com's LinkBuilder® Multi Services Hub. The most advanced, cost-effective networking solution in the industry.

Thanks to its unique "technology-independent" backplane, rather than adding backplanes for every new technology, the LinkBuilder MSH™ allows you to mix Ethernet and Token Ring networks today, and add ATM and other new networking technologies when you need them. All the while avoiding costly and disruptive hardware upgrades.

But why not see for yourself. Just call us a call at **(02) 959-3020, ext. 20** for a free copy of 3Com's white paper "High-Performance Scalable Networking," and more information about the LinkBuilder MSH. You'll learn the easiest way to plan for the future is to have a hub that's already there.

3Com®

Networks That Go the Distance™

Anixter Australia	(02) 333 0800	Bull HN Information Systems	(02) 847 7100	Fujitsu Australia	(02) 410 4555
GEC IT	(02) 415 0500	Interconnect NZ	(649) 358 0322	LANCON	(02) 899 3033
NetArch	(07) 393 1933	NETSOURCE	(09) 470 5311	NST	(02) 499 2900
Scitec	(02) 428 9555	Tech Pacific Australia	(02) 697 8666	Tech Pacific NZ	(649) 415 8300

nologies that vendors could quickly adopt to bring products to market.

At the physical layer, HPPI is defined by two standards. HPPI-PH (HPPI-physical) covers copper-based parallel cabling that boasts a whopping 800Mbps data rate. If that is not enough bandwidth, the HPPI standard provides for extending it to 1.6-Gbps by doubling the width of the data path. However, use of copper rather than fibre to obtain these substantial data rates also implies rather strict distance limitations — a maximum of about 25 metres. The other part of the HPPI physical standard is HPPI-SC (HPPI-switch control), which defines switch control that translates the point-to-point nature of copper cabling into the any-to-any connectivity characteristic of other emerging channels.

Data Framing

The middle layer of the HPPI spec is HPPI-FP (HPPI-framing protocol), which defines data framing. This portion of the specification is designed to accommodate both large block transfers and framing to split data into smaller bursts. Also defined under FP are identifiers by which HPPI channels can support multiple upper-layer protocols.

Upper-layer standards include HPPI-LE (link encapsulation), which maps logical link control (LLC) frames (IEEE 802.2) onto HPPI for communication between pairs of hosts. HPPI/IPI-3 maps IPI-3 (Intelligent Peripheral Interface) command sets, the same ones that run over Fibre Channel, over HPPI physical layers for communication between hosts and peripherals.

HPPI channels enjoyed a fast start in the late 1980s and early 1990s, but the future is a question mark. On one hand, HPPI's acceptance, particularly by the scientific and high-performance market, has given it a head start over the other emerging channel technologies. In addition, its 800Mbps data rate exceeds available implementations of most other channel technologies. Best of all, it is available today using reliable technology. Conversely, Escon has only recently become available and Fibre Channel is just getting off the drawing board.

On the other hand, certain elements of HPPI appear to pose problems that may affect its future. For example, although many high-performance hosts are available with HPPI channels, few peripheral vendors have delivered HPPI/IPI-3 products. This means that one of the central roles of channels, communication between hosts and peripherals, will remain impractical until more IPI-3 peripherals become available. In contrast, Escon has received support from virtually all of the peripheral vendors in the IBM mainframe marketplace, and Fibre Channel seems to be gaining momentum in the scientific, high-performance, and even the commercial marketplaces. Although few expect that this factor will cause HPPI to disappear in the near future, some fear that without greater use of IPI-3 the future role of the HPPI channel may be limited.

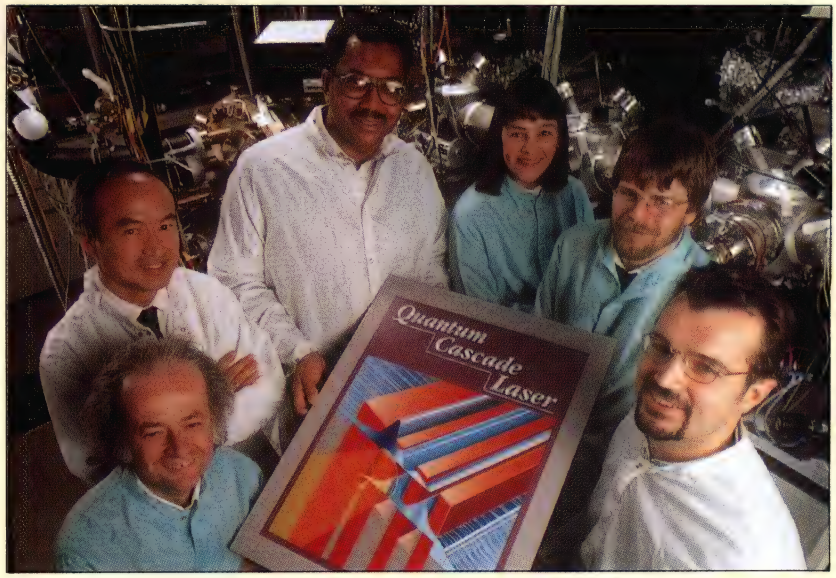

AT&T Bell Labs Laser Breakthrough

Scientists at AT&T Bell Laboratories have invented a fundamentally new kind of semiconductor laser which operates like an electronic waterfall. The new quantum cascade (QC) laser is the first laser which can be tailored to emit light at a specific wavelength which can be set at nearly any point over a wide range from the mid- to far-infrared spectrum. When an electric current flows through a QC laser, electrons cascade down an 'energy staircase,' emitting an infrared photon, or light pulse, every time they hit a 'step.' At each step, the electrons make a quantum jump between well-defined energy levels, and the emitted photons are reflected back and forth between built-in mirrors, stimulating other quantum jumps and the emission of other photons until the amplified pulse escapes the laser cavity. The new laser, the culmination of 30-years' research, was developed using band-structure engineering and molecular beam epitaxy technology. It has potential uses in a broad range of applications, including monitoring of air quality, industrial process control, free-space point-to-point communications, and spectroscopy.

For the LAN

Just as channel speeds are increasing, emerging LAN technologies are offering improved throughput over established techniques. Although not usually thought of in terms of being a channel, these technologies nevertheless begin to approach the definition of a channel in many respects. For example, LANs can be used for one of the major channel applications: host-to-host communication.

Newer LAN technologies also boast data rates that typically exceed the speed of heritage channels and rival those of the emerging channels. For example, FDDI and 100Mbps Ethernet put LAN transfer rates on a par with Escon.

Like Escon and Fibre Channel, FDDI uses fibre optics as the transmission medium and therefore has a range of several kilometers. In addition, both channels and the emerging LANs are designed with a high degree of resilience in mind. But FDDI offers an added degree of redundancy thanks to its dual rings. Should one of the fibres break, the devices on the ring automatically perform a 'ring warp' that changes the flow of information to bypass the failure. Alternatively, dual-fibre FDDI can be used to double network capacity. These factors are positioning FDDI as a high capacity, resilient LAN.

One drawback to FDDI is that, in contrast with the switched-channel technologies, it uses a shared-media approach that, despite its dramatically increased bandwidth, tends to effectively limit available bandwidth. This particularly true for networks that involve multiple channel connections to the LAN.

ATM on Arrival

Another technology often thought of in terms relative to channels is ATM (asynchronous transfer mode). Although its primary purpose is to serve as a combination of a LAN fabric and a WAN fabric, it also features many properties of channels, leading to the question of whether it will ultimately take on this role as well.

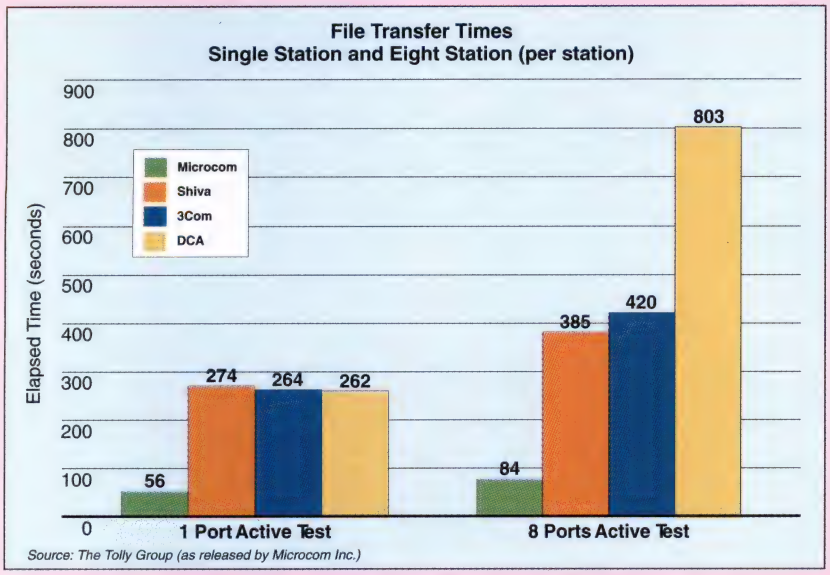

Source: The Tolly Group (as released by Microcom Inc.)

Microcom Claims Top Marks with LANexpress

Tests commissioned by leading modem vendor, Microcom, and carried out by independent test lab, The Tolly Group, indicate Microcom's remote LAN access system, LANexpress, can deliver greater data throughput than competing products. The tests measured the file transfer performance of four dial-up LAN access servers — 3Com's AccessBuilder (formerly Centrum Remote), DCA's Remote LAN Node, and Shiva's LANRover/8E — in an environment where multiple remote users are simultaneously accessing the LAN via a single LAN access device. The tests revealed LANexpress can achieve aggregate throughput of 800,000bps, beating competitors by up to a factor of ten. In multi-port tests, LANexpress also delivered 100,000bps simultaneously over all eight server ports tested, while other products suffered performance degradation. Commenting on the tests, Tolly Group President, Kevin Tolly, said "LANexpress was the only product not plagued by significant operational problems. The multi-port test clearly demonstrated that LANexpress has ample scalability when network activity increases dramatically."

Again, data rate is one of the most obvious parallels between ATM and channels. Currently available speeds for ATM range from OC-1 (DS-3 speed or 44.736Mbps) to OC-3 (155.2Mbps), and defined speeds run all the way up to OC-48 (2.5Gbps) or even more. With that much bandwidth in the ATM specification, it is capable of taking on just about any role it wants, including that of a channel.

Another key characteristic of ATM is a fixed-length, 53-byte cell format intended to provide a predictable delay and a simpler formatting process. This also makes ATM suitable for a wide variety of transmissions including voice, video, and data.

In addition, ATM holds the promise of providing a seamless 'on ramp' to high-speed, wide-area ATM services. Since ATM as it is defined in the local area is the same specification that carriers are using to develop long-haul services, moving ATM cells from a LAN across a WAN becomes completely transparent. In the context of a channel, this would give ATM a key advantage over other channel technologies in that it would effectively extend the length of the channel over unlimited distances.

The 53-Byte Limit

But ATM's 53-byte cells, although well suited for the smaller packets typical of LAN traffic, voice, and video, are relatively small compared with the 4K to 32K blocks typical of mainframe traffic. This means that some device, either the channel interface itself or an external unit, will need to slice and dice the blocks of channel data into cells at the originating node, then reassemble them at the remote end. Further, each cell dedicates a relatively inefficient 5 bytes to overhead, which will be multiplied hundreds or thousands of times as channel blocks are mapped into ATM cells.

This situation causes a unique problem for ATM as a channel technology. The scheme currently built into ATM to resolve congestion is to simply drop cells; since ATM was originally planned as a voice-oriented technology, dropping cells would have little effect on the parties involved. However, cells dropped in a data transmission ruin the entire transmission, which means the sender must start the communication again. In the context of host channels, this would mean resending the entire block, not just the cells lost during the transmission. The ATM Forum is actively debating various approaches to this issue, primarily by employing techniques that reduce the instance of dropped cells.

This becomes a particular challenge for ATM since other channel technologies already address this problem. For example, another key characteristic of Fibre Channel is in the way the standard handles flow control. According to the specification, Fibre Channel allows the receiver to specify the amount of resources it has available at any given moment, guaranteeing that the sender does not transmit too much information. This eliminates congestion on the network that would force ATM to drop data, and guarantees that data sent through Fibre Channel can be handled by both the network and the destination device.

However, ATM does offer some unique strengths, particularly when extended distances are involved. Remember that even the emerging channels suffer from distance constraints, and one of ATM's characteristics is its seamless transfer across WAN services. In addition, channels have always gravitated toward the WAN services offering the greatest available bandwidth when needing to overcome their distance constraints, and ATM certainly fits that description. The upshot is that ATM is considerably more attractive as a WAN for channels than it is for a channel interface itself.

Other Means

However, ATM is not the only means of moving channel traffic across a WAN; in fact, most channel networking technologies today make substantial use of high-speed circuits for just this purpose. In these applications, it is important to keep in mind that substantial cost savings can be realised by combining multiple applications across a single WAN. If there are channels on both sides of a WAN circuit and LANs on both sides of a circuit, why not leverage the economies of scale offered by higher bandwidth circuits for both applications?

Both channels and LANs may be connected to a single WAN in a variety of ways. Several classes of devices, including routers and channel networking systems from CNT, Network Systems Corporation and others offer direct connections to WANs. These devices can often accept data from other sources and tunnel this traffic travels across the WAN. Alternatively, various bandwidth management products such as those from US vendors Adaptive Corporation, Digital Link Corporation, and T3plus Networking furnish a similar function.

These products accept data from multiple sources by using either of two standard interfaces: V.35 (up to 4Mbps) or HSSI (high-speed serial interface, 52Mbps). The products then multiplex this traffic onto a single WAN circuit.

Technology Update

Ethernet to Get Speed Sensors

Network managers connecting new Ethernet equipment that can work at either 10Mbps or 100Mbps are presently faced with the somewhat daunting task of manually setting all hub ports and adaptors before connecting them. To solve this problem, the technical team at National Semiconductor has been hard at work developing the company's new Nway software, which enables multispeed Ethernet equipment to sense the link speed and automatically adjust to it. When hardware is powered up, Nway's protocol automatically negotiates a connection at the highest speed supported by the devices at either end of the link. Officials said the software can also be used to help equipment self-configure to operate on 20Mbps and 200Mbps full duplex connections, or on 16Mbps Iso-Enet LANs. The Nway software has already been approved by the IEEE 802.3 Working Group, and may be included as an optional addendum to the 802.3 Ethernet standard.

Wireless X.400 Standard

The vendor consortium known as the Cellular Digital Packet Data Forum has begun work on a new X.400 wireless standard which will allow e-mail to be sent over CDPD (Cellular Digital Packet Data) networks. Two fundamental issues need to be resolved in the spec — how to work around the different packet sizes of Token Ring, Ethernet and other network technologies, and how delays will be handled so that sessions don't time out. The spec will eventually be submitted to the US Telecommunication Industry Association (TIA) for formal approval. Meanwhile, Motorola has teamed with router vendor Retix to develop a problem-free wireless billing system. Retix will provide X.400 technology for Motorola's Celtac CDPD system that will let carriers provide users with a consolidated bill, overcoming accounting problems caused when CDPD users roam from their home territory.

Rockwell Unveils World-Fastest DAC

Rockwell Telecommunications has announced what it claims is the world's fastest Digital-to-Analogue Converter (DAC). The product converts digital data from fibre cable to analogue signals that can be received by existing television sets, potentially giving users access to hundreds of TV channels as well as voice, data and multimedia capabilities. The 10-bit DAC, which has been designed and manufactured by Rockwell's Microelectronics Technology Centre, operates at clock and data speeds of more than 1200MHz, while consuming less than 800mW of power. It is based on Heterojunction Bipolar Transistor (HBT) technology and has an extremely fast settling time — less than 1nS to 1/2 Least Significant bit, with a glitch impulse of less than 1pV-s. It has industry-standard Emitter Coupled Logic (ECL) interfaces, and is designed to operate on a single power supply of -5.2V. Officials said the DAC can also be used in Direct Digital Synthesis, where, coupled with a digital Sine-ROM accumulator, it replaces phase locked loop-based synthesizer circuitry in digital radio frequency transceivers, potentially reducing the cost of spread spectrum systems.

HP's Coup: AT&T Adopts OpenView

Champagne corks must have been popping at Hewlett-Packard's OpenView development facility in Colorado recently with the signing of a deal to license the company's network and systems management platform and applications to AT&T. The deal makes HP OpenView the world's most widely adopted standards-based network management platform, according to HP officials. AT&T will use OpenView, which now incorporates telecommunications management, as the basis of its new company-wide OneVision Network Management Solution, which will manage the company's internal network, and which will also be offered as a solution to AT&T customers. Over 150 vendors have now licensed the OpenView technology, including Alcatel, Siemens, Data General, Groupe Bull and Hitachi, and HP says the software is currently installed in more than 35,000 networks worldwide.

SPA Claims Codec Breakthrough

Melbourne-based firm Signal Processing Associates (SPA) has developed a codec which provides voice quality approaching that of 64-Kbps encoded voice at the low code rate of 16Kbps. Hailed as a breakthrough, SPA developed the implementation of the software algorithm for the codec under a technical development contract with Telstra's International Engineering group. Telstra said the codec has already drawn interest from telecommunications equipment manufacturers around the world. The algorithm was initially developed by Telstra for use in satellite modems in its global digital DAMA-Net service, and will now be commercialised for a range of applications including videophone, store and forward fax, ISDN and multimedia.

Intel Debuts Smooth Motion Video

Intel has announced version 3.2 of its Indeo PC video technology, which the company says delivers a dramatic improvement in picture quality, size and playback performance for all Pentium and 486-based systems. Officials said the Indeo software-based codec's scalability enables it to deliver smooth motion playback of a video file in a 320 by 240-pixel window on a 486, or full-screen on a Pentium-based system. The addition of a DCI-enabled graphics card to an appropriately-configured Pentium system allows the video clip to be played full-screen at VHS quality. Apple has recently revealed it will incorporate Indeo technology in QuickTime 2.0 for both Windows and Mac.

OSI-IP Backers Seek Common Ground

EDICA's recent forum in Melbourne seemed to find more hostility than common ground between supporters of OSI protocols and the Internet Protocol (IP). The OSI camp touted the OSI suite's open architecture, security and audit functions, and high-level system management tools, while IP proponents countered that the high cost of implementing OSI effectively made it an option for only the largest organisations. Common sense prevailed, however, in the form of Network Manager of the Australian Stock Exchange, Mike Aikins, who said the debate should focus on the real business needs of users, pointing out that both technologies already borrow from each other — for example, the Internet's use of X.500 directory services, and OSI applications which make use of Internet protocols. A further meeting to draft an industry 'co-operation plan' was scheduled for late July.

Video Over Spread Spectrum Wireless

Australian wireless communications specialist, Pinnacle Communications, has trialled a system for transmitting video images from a helicopter back to a land-based system using a pair of Cylink Airlink 64Kbps spread spectrum radio modems. The system, which was developed for Esso's Hastings plant on Victoria's Western Port Bay, used an Adpro fast scan colour video transmitter/receiver, and successfully transmitted good quality survey images from the helicopter back to the base 13 kilometres away. Officials said that with minor antenna engineering the system should be able to cover up to 20 kilometres.

Philips Teams With US Partners

Philips' US operation has signed agreements with TV company Zenith Electronics and videoconferencing developer Compression Labs to develop and market TV set-top decoders for digital and analogue networks which will provide two-way interactive communication to services such as video-on-demand. The units will incorporate MPEG-2 digital decoder technology from Philips and Compression Labs, and Vestigial Sideband transmission and real-time two-way technologies from Zenith. The digital decoders will be compatible with video servers under development by other vendors as well as with Compression Labs' MPEG-2 encoders and Zenith's head-end equipment, including analogue scrambling systems and digital VSB modulators.

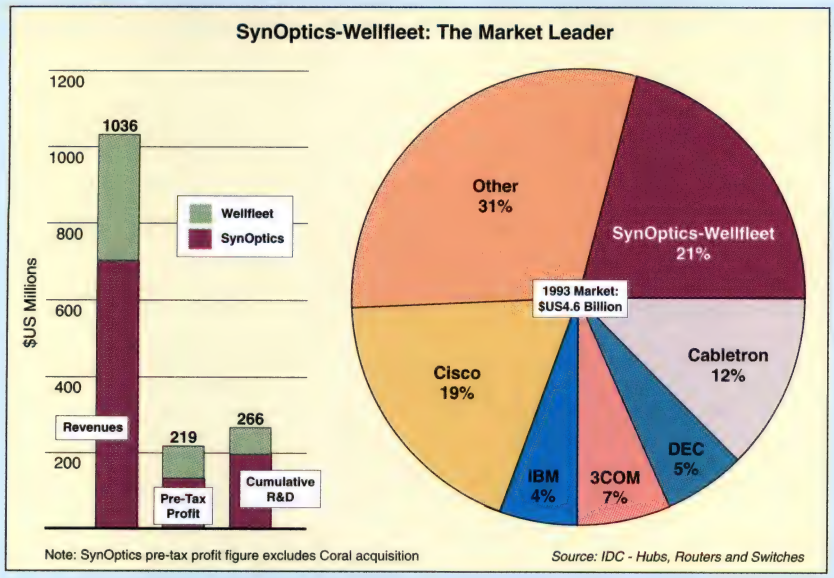

Wellfleet, SynOptics Merge For Market Leadership

Market-leading hub vendor, SynOptics Communications, and high-end router manufacturer, Wellfleet Communications, surprised the networking industry last month by announcing plans to merge. The move, which caught most US industry analysts off-guard, will form a new global company which was still unnamed at time of going to press, and which will hold some 21% of the internetworking product market according to market researchers, IDC. While logistical details remain fuzzy and both companies are vague about the new corporate structure, it is believed Wellfleet's Billerica, Massachusetts, headquarters and SynOptics' home in Santa Clara will both be retained, and that the companies will continue to pursue existing alliances with other vendors. Both companies already share important partnerships with Novell, Intel, and US telco WilTel. SynOptics, which claims to have shipped over 10 million network nodes, also has strategic relationships with IBM, AT&T, Sun and Sprint. Wellfleet, twice-named as the fastest growing company in the US by *Fortune* magazine, has in place alliances with Banyan, and, perhaps more importantly, Microsoft. This latter relationship has been strengthened recently with the signing of a multi-faceted agreement covering joint product development, testing, support and sales and marketing of internetworking products, including the Windows NT operating system.

The announcement of the SynOptics/Wellfleet merger met with mixed responses from industry analysts and competitors, with most admitting the new entity has the potential to be a very strong force in the internetworking market, but with some expressing doubts that the two companies will be able to successfully blend their very different corporate cultures and marketing approaches. Meanwhile, market-leading router vendor Cisco has announced the acquisition of Newport Systems Solutions, a niche router vendor.

One approach to sharing WAN circuits is to simply assign a predetermined slice of the bandwidth available on the circuit to each application. This ensures that each device gets access whenever it needs it, but since many applications involve substantial idle time, it also has the effect of limiting throughput when other applications may be starved for bandwidth.

Another approach to bandwidth management is to prioritise traffic. In this scenario, high-priority applications get the full access to the bandwidth available on the circuit, while lower-priority applications get buffered and transmitted only when the circuit is available. Although the lower-priority applications wait for the circuit to become available, they too get full access to the rated capacity of the link when their priority slot becomes available. The net effect usually is to improve throughput levels for all applications.

Channel to Channel

In all of this, note that emerging channels represent little more than improved physical transports for existing logical command sets that have been around for some time. This prompts a series of questions relating to if, how, and when the players in the market will take the next step: will an additional degree of interoperability be established by providing a consistent command set and blocking philosophy across physical channels?

One subject of speculation is whether or not IBM will develop a Fibre Channel interface for mainframes as the successor to Escon. One negative sign: Escon has been in place for some years, and IBM has made no public statement concerning such a policy while busily rolling out additional devices under the Escon umbrella.

Nevertheless, arguments in favour of IBM's taking such a step appear to be strong. With the development of Fibre Channel on the RS-6000, it makes a certain amount of sense for IBM ensure interoperability between members of different product groups. IBM's substantial presence in the Fibre Channel Committee suggests a vested interest: in fact, IBM has played a substantial role in defining many of the Fibre Channel specifications.

In addition, IBM and Amdahl have worked together to develop a Fibre Channel FC-4 called FC-SB (single byte), which in many respects is a fancy phrase for the CCW structure used on bus-and-tag and Escon channels. This standard parallels the concepts embodied in Escon, implying that an intention (or at least a by-product) is to make the migration from Escon to Fibre Channel a relatively painless task. The same issue applies to SCSI. Actually, even today some vendors are running the SCSI command set over Fibre Channel. For example, Sun Microsystems recently unveiled a disk array that works in exactly this manner.

With communications technology changing at such a rapid pace, it is no wonder that network designers sometimes struggle with questions relating to how much bandwidth is required or which standard to select. The explosion of bandwidth associated with the emerging channel technologies has the potential to alleviate this quandary, particularly since it has been matched by an equal explosion at the LAN and WAN levels.

But the migration from established technologies to the new generation of high-bandwidth, fibre-based technologies represents more than bandwidth alone. Equally significant is that these technologies are melding channels, LANs, and WANs in ways that lower or eliminate some of the most pronounced barriers between computing platforms today. The consistent application of fibre optics and the flexibility with which the emerging channels handle protocols from alternative platforms suggests that channel technologies are converging in ways that will extend their opportunities for connectivity.

The fact that the emerging channels embody a switching architecture makes the channel a network of its own, and opens the door for using the channel as an integral part of the enterprise-wide network.

Jim Morin is the Director of systems marketing Computer Network Technology, a US-based manufacturer of channel networking systems. Gene Misukanis is the vice president of engineering at CNT.

Value.

Is not only
what you pay.
It's what you get
for your money.

- ✓ Ethernet/10BaseT
- ✓ Token Ring
- ✓ SNMP Management
- ✓ HP-OV Software
- ✓ Fibre Modules
- ✓ Built-in Port Redundancy
- ✓ Automatic Recovery Function
- ✓ UTP and STP Cable Support
- ✓ 16-Ports
- ✓ Simple Upgrade
- ✓ Independent Sub-Networks

Quality Endorsed Company
AS 3902 Lic 3012
Standards Australia

All of these features are now available in ONE stack with the new INC 516 series of Multi-protocol Stackable Hubs.

With INC's commitment to quality products at sensible prices, all these features are now available at a similar price per port to our 412 Series 10BaseT hubs which have been so well accepted in the marketplace.

Call INC Manufacturing Pty Limited on 008 807 982 or MPA on 008 801 767

INC

I.N.C. Manufacturing Pty Limited ACN 051 977 701
Unit 2, Block R, 391 Park Road, Regents Park NSW 2143

Wireless Communications

PCMCIA Puts Paging Into Portable PCs

Remote communications will soon get easier, and perhaps cheaper, as two best-selling technologies are merged together through the use of PCMCIA (Personal Computer Memory Card International Association) cards. Portable PCs and alphanumeric paging will soon go hand in hand, in the new round of hardware and software products set to hit the streets.

According to US researcher BIS Strategic Decisions, portable PCs will become the best-selling mobile device by 1997, with pager sales close behind with an average growth rate of 10.6% a year.

The Personal Communications Industry Association (PCIA), has also predicted rapid growth for wireless services over the next ten years, with paging services leading this market, even over cellular phones.

Merging the two technologies will not only let the user take advantage of simpler and less expensive means of sending information, but offer pager manufacturers an 'in' to a very fast-growing markets. The lower power requirements also make paging conducive to the portable PCs, allowing the devices to be of a lighter weight.

One of the first of this new breed of mobile communications products has been developed by US vendor NovaLink. The company has released a PCMCIA Type II card, the NovaPak 144, that combines not only a wireless alphanumeric receiver but a 14,400bps data, fax and voice modem. The NovaPak 144's voice messaging capabilities also allow for multiple mail boxes and private outgoing messages.

Toshiba has also announced a PCMCIA type II receiver card for notebooks. The company's Noteworthy NewsCard allows mobile users to receive pages and receive and display wireless e-mail messages.

While for many years numeric paging has meant the ability to send a short text message to another person — essentially 'notification' — the next generation of alphanumeric pagers will instead comprise wireless messaging centres providing actionable information which mobile professionals can do something with immediately.

And unlike cellular, alphanumeric paging/wireless messaging systems permits one-to-one as well as one-to-many applications. This can be an added benefit in circumstances where a company needs to communicate the same message to many staff, such as a mobile sales force. Instead of having to call each rep, as with a cellular phone, alphanumeric paging can deliver the message to everyone at the same time.

Monica Snell

Intelligent Networks

Network Problem? Call My Agent

The telecoms industry will soon become familiar with the term 'intelligent agents,' a new software model which is beginning to penetrate the world of networking and communications, and which may become the most significant advance in software technology to hit the market since the GUI.

Described as 'objects that think,' an agent is a piece of software code which (often using artificial intelligence) executes part of a programmatic process. Agents are already now being embedded for use in network management, but will also begin to appear in message management and information retrieval software.

Companies such as SynOptics, Alcatel and BT are already placing 'embedded intelligence' into their network products, with other vendors expected to follow suit. SynOptics in particular is already well advanced in the development of agent technology for its global enterprise management (GEM) architecture, which is intended to eventually manage all corporate communications within a single application. The company's 'SuperAgents' operate at a distributed management level and collect data from network devices, analyse it and take the appropriate action. SynOptics says its ultimate aim is to develop self-healing and self-managing networks.

Besides network management, agents are also expected to become an integral part of messaging systems. AT&T's PersonaLink messaging service, scheduled for debut this year, uses agent technology — notably, General Magic's Telescript — which supports electronic 'intelligent assistants' using object oriented code. Bellcore, IBM and Apple are also using agents in new product development — Apple specifically for its Newton and System Software, and IBM for its new Rules technology, which will utilise aspects of agent technology to allow products to accept input from equipment including telephones, fax machines, server applications and internal e-mail.

Behind the golden promise of these new agents, however, lurk some serious security risks. One danger is that of specialised assassin agents created by hackers which may enter a network and cause damage or 'kill' other agents. Another area of risk is controlling agents which learn from other agents and take action based on that knowledge — action that the user might not want taken.

Martin Cheek

Enjoy the cost savings, revel in the performance

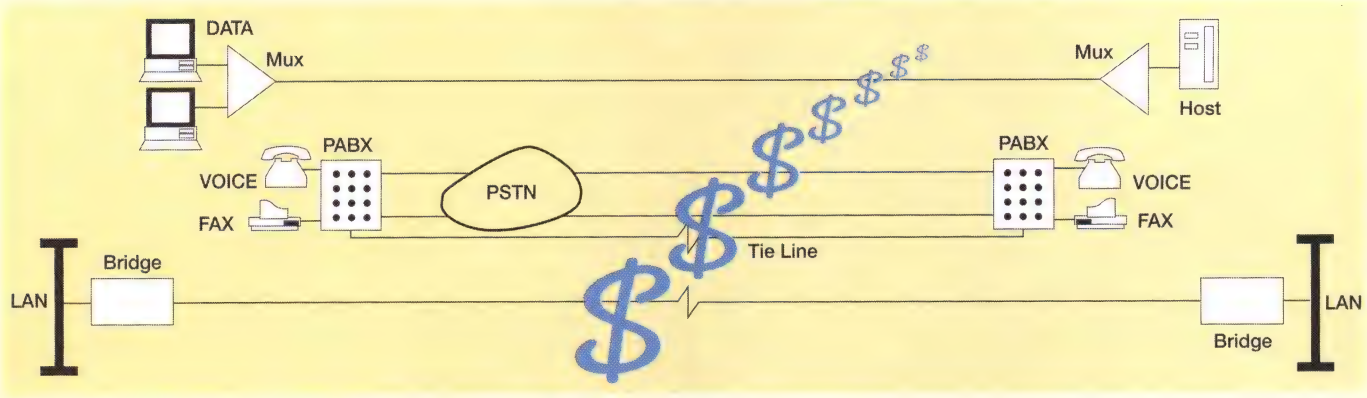

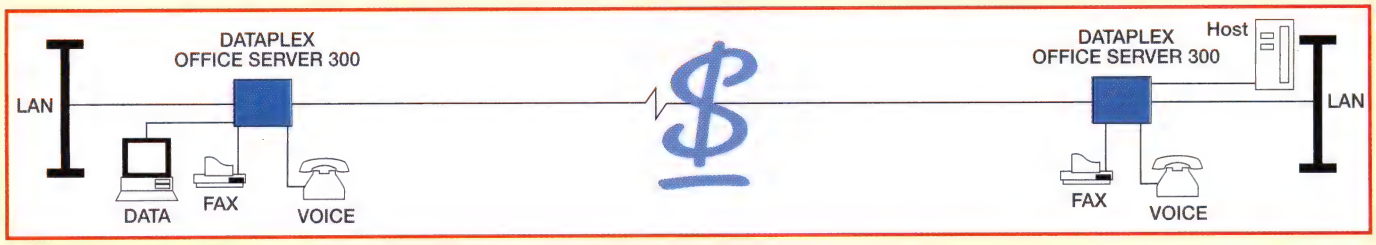

of the Dataplex Office Server 300

Pay-back periods of less than 12 months.

On-going line costs reduced by 300%. . . and more.

Voice, data, fax and LAN traffic on the one line instead of multiple lines.

Savings like these, and dramatically improved network performance are possible with the Dataplex Office Server 300.

Using Fast Packet Multiplexing which allocates bandwidth based on the requirements of data, voice, fax and LAN traffic, and the latest compression techniques, you obtain maximum network performance while using the minimum bandwidth.

Benefits of the Dataplex Office Server 300:

- Improve your inter-office Voice, Fax, Data and LAN communications
- Increase the capacity and flexibility of your communications network
- Minimise your existing and future capital investment
- Maximise the investment in your existing networking equipment
- Simplify your network management and maintenance costs.

The Dataplex Office Server 300 has Austel approval and offers a very cost effective solution for integrated data, voice, fax, and LAN networking.

Call Dataplex today. We'll send you more information.

Dataplex

MELBOURNE TEL: 03 210 3333 FAX: 03 210 3399
SYDNEY TEL: 02 349 3911 FAX: 02 349 7747
CANBERRA TEL: 06 285 3657 FAX: 06 285 3696

DPX 4

Uncovering network problems one segment at a time could leave you open to some unpleasant surprises.

HP NetMetrix distributed monitoring solutions: Now nothing can hide from your view.

Here you go again. There's a hot spot in your network and you don't know what's causing it. So you begin searching. Segment, by segment, by segment. Meanwhile, the source of your problem continues to raise havoc.

What can you do about it?

Put an end to your search with HP NetMetrix. Our distributed monitoring and analysis solutions allow you, from a single computer console, to instantly know what every segment of your network is doing – whether it's across the hall or across the world.

Take our correlated viewing capability. It shows you how activities on different segments affect each other. So you can quickly find problems based on insight, not intuition.

Of course, your goal is to catch problems before they become a disaster. With HP NetMetrix, you can set networkwide alarms to alert you at the first sign of trouble.

Flexibility and compatibility are other HP advantages. All our solutions are open and are based on the SNMP and RMON standards. They integrate with all leading network management platforms, such as HP OpenView and SunNet Manager, and support a variety of computer hardware platforms.

Best of all, because having a complete view helps you better understand your network, you can make more informed architectural decisions and optimise your network's performance.

Want proof about how HP NetMetrix can help you uncover problems and manage your network like never before? Call our Customer Information Centre on **13 1347** (Australia wide) and ask for extension **2902** for literature on our distributed monitoring solutions.

You'll never look at your network the same again.

A better way.

HEWLETT® PACKARD

SOFTWARE VIEWS

IBM NetView/6000: No Match for HP OpenView

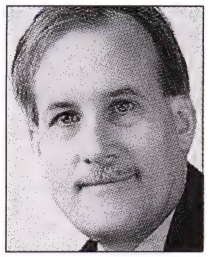

Frank Henderson

To hear IBM tell it, Big Blue took a look at HP OpenView 4.X and decided not to renew its source code licence with Hewlett-Packard. It seems as if HP's next-generation network management platform, which boasts an object-oriented architecture, isn't technologically advanced enough. I guess IBM was equally unimpressed with the fact that HP OpenView 4.X will allow large corporations to geographically distribute management services — including configuration, event handling, security, and accounting — across the enterprise.

For its part, HP is only saying that it couldn't reach an agreement with IBM. But I bet there are a lot of smiling faces out in Fort Collins, Colorado, where the OpenView development team is located. And I think HP is going to have the last laugh. Pretty much any way you slice it, HP is way out in front of IBM when it comes to net management. Big Blue obviously believes that it can bring NetView/6000 up to speed without HP's help. I think a quick reality check shows how wrong IBM really is. Maybe IBM has started to listen to the self-proclaimed industry experts who never tire of telling us that NetView/6000 is the No. 1 choice for big companies looking to manage client/server networks. Too bad the numbers don't bear out those claims. According to the 1993 survey of Unix SNMP management consoles published by International Data Corp. (IDC), Sunnet Manager from Sun Microsystems has 30% of the market; OpenView from Hewlett-Packard, 26.7%; and IBM, 6.6%.

But there's more to this than market share. IBM originally licensed some OpenView modules in 1991, with the agreement running all the way through OpenView 3.X. NetView/6000, in fact, could be accurately described as IBM's version of OpenView. But IBM never bothered to lock in the APIs (application program interfaces) that third-party ISVs (independent software vendors) use to integrate their applications with either NetView or OpenView. Now that the HP-IBM deal is a no-go, Hewlett-Packard is free to rewrite and optimise its APIs. If it chooses to do so, any applications written to the NetView/6000 APIs will run less efficiently under OpenView — if they run at all.

"Pretty much any way you slice it, HP is way out in front of IBM when it comes to net management."

One of the key elements of version 3.X of the HP platform is OpenView Network Node Manager, which automatically discovers IP and SNMP nodes and builds the network map. Node Manager also delivers dynamic data collection to correlate MIB (Management Information Base) data from devices on the network. HP last year announced its OpenView Operations Centre, an event-handling application that takes care of system management tasks like backup, server performance monitoring, software distribution, terminal management, and the like — all from the OpenView interface. What's key here is that Operations Centre is a systems management application that has been integrated into the OpenView network management platform, thus furnishing excellent efficiencies of scale for overseeing and troubleshooting system and network nodes.

HP also delivers systems management with its Extensible SNMP agent, a mechanism for managing Unix workstations via a standards-based management platform. Extensible mapping allows SNMP management programs like Network Node Manager to perform customised Unix system administration.

NetView/6000 offers the same basic components as HP OpenView Node Manager. IBM has added some features of its own; all in all it's done a good job simplifying the interface to its system and creating administrative tools. For example, the System Management Interface Tool (SMIT) offers menus and prompts to help with installation and network management.

But IBM has done very little to enhance the essential functions that it has licensed from HP. Almost its only venture on this front is the General Topology Manager (GTM), an application that enables information about any node to be automatically displayed on the network map. It's a nice program; too bad most third parties have ignored it. Even IBM partner Chipcom has shown no interest.

IBM also has released AIX Systems Monitor/6000, which makes it possible for the standard RS/6000 SNMP agent to capture systems management information. The extension automatically filters the amount of data sent to the central management console, cutting down on unnecessary traffic. But Systems Monitor/6000 is pretty rudimentary when compared with HP Operations Centre. Worse, it's not fully integrated and needs lots of fine-tuning.

HP has also put RMON (remote monitoring) MIB to work with OpenView. This is critical, since SNMP on its own gathers and passes along much more information than one central console can keep track of. The HP Probe Manager captures RMON data and displays it to the OpenView interface, using the same graphing utilities for both on-line and historical reports. At this point, IBM has no generally available RMON-integrated applications. Protocols did deliver beta RMON code for OS/2 and promised to do the same for AIX. Now that Network General has brought out Protools, though, the outlook for both projects is bleak. Recently, IBM and Technically Elite Concepts announced an RMON/6000 application. When it will be truly integrated with NetView/600 is anybody's guess.

But technology is not the entire story; To be truly effective, a net management platform must be backed by a comprehensive support structure. HP has come up with its Premier Partners program for ISVs and other divisions to certify that software will perform as promised; develop technology and process management; and certify systems integrators who can install and customise OpenView. IBM is also trying to establish a forum for vendors and end users. The NetView/6000 Association attempts to update application availability on the platform and certify software performance.

On paper, the two sound very similar. How do they really stack up? Leading-edge US systems integrator Applications+Plus is the only charter member of both groups. Clayton Rote, the company's senior marketing manager, sums it up this way: "IBM appears less committed than HP. The HP Premier Partners is truly an open forum that allows ISVs and systems integrators to drive releases in upcoming OpenView versions."

So if you hear HP chuckling, you'll know why. Or should that be laughing all the way to the bank?

Frank Henderson is Chief Technology Officer of the The Netplex Group, a systems integration firm based in Maplewood, New Jersey.

LAN VIEWS

The Use and Abuse of Synchronous Access Schemes

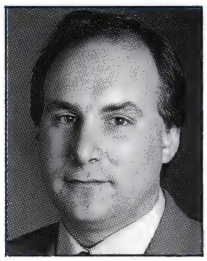

Graeme Le Roux

With its low cost, ease of management, simplicity and reliability, 10Base-T Ethernet has become the undisputed first choice (at least this year) in desktop networking technology. But while it may have fitted requirements so far, the client/server application revolution is beginning to highlight 10-Base-T's shortcomings, and some network managers are discovering that an asynchronous access scheme like Ethernet is no longer the best backbone choice. The reason for this is that the latency of an asynchronous access scheme depends upon instantaneous load and the dependence is not linear. This is precisely the sort of environment you want to avoid for applications which perform synchronous operations such as 'multi-phase commits.'

A multi-phase commit is an operation where a record in a server-hosted master database is modified by a client application and, for security purposes, a remotely-located copy of the database on another server is updated at the same time. The details of this mechanism vary with implementations, however the main point is that unless both servers report completion of the modification within a time period built into the application, the transaction fails, an error is reported to the client and any modifications to the record which may have been made by either server are 'rolled back' — i.e. the record is guaranteed to be in its original state on both servers.

From the point of view of network design, the details of the application's actual task is of little interest; what is critically important is the fact that the time delay (latency) between the application client and the application servers be predictable and that the worst case latency in the path over which any time-critical operation is to be performed will not exceed any time-out value for that operation which has been built into the application. In order to be able to meet this requirement a network designer must understand the way in which the application clients and servers communicate — as opposed to what they are communicating.

To return to the multi-phase commit example, there are generally two ways in which this operation will be performed. Either the application client will write to both servers and wait for both servers to confirm the transaction, or the client will write to the master server which will then, transparently, replicate the operation to one or more backup servers and then confirm the transaction to the client once all servers have confirmed their individual updates. In the former case, it is the latency between the client and the servers which must be predictable, while in the latter case it is the server-to-server latency which is most important. Fortunately for network designers, it seems that this latter case is the most common, particularly at large sites, because it makes guaranteeing the range of variation in a network path's latency reasonably straightforward — you use a synchronous access scheme like Token Ring or FDDI (or C/SDDI).

Synchronous access schemes were originally developed for inter-host communication (what we now refer to as a system backbone) and were designed to eliminate the possibility of a system becoming overloaded and thus its station-to-station latency varying widely. When IBM started marketing the original Token Ring LANs against Ethernet — which was then coax-based — one of their main selling points was superior throughput under load. This was typically demonstrated by having a number of stations copy lots of files around. They also pointed out that Token Ring could detect some media faults and had rudimentary management capability. The other thing which Token Ring has over Ethernet is a larger frame size.

Unfortunately Token Ring is more expensive to deploy on the desktop than Ethernet and does not provide some of the features found in 10Base-T implementations. Most synchronous access schemes are also difficult (and therefore expensive to deploy across multiple bridges), and do not handle floods of small packets very well. But as a private backbone for servers or hosts which are running distributed applications synchronous access schemes are, technically speaking, the best possible choice since their capacity to handle heavy loads of large packets between a reasonably stable number of stations reliably means that such applications can be designed on the basis of a network path latency which is at best invariant (fixed number of stations) and at worst deterministic. Schemes such as FDDI can also be configured to be fault tolerant far more easily than an Ethernet-based system.

Nor are synchronous access schemes much more expensive than a fault tolerant Ethernet solution. If you have several servers in a single location which run one or more distributed applications then all that is required is a single MAU/hub, appropriate adaptors and some cables. Even CDDI, which uses Category 5 UTP cable, is often justifiable provided one forgoes the fault tolerance provided by Dual Attached Stations (DAS). Given that all your servers are locked away in a secure and stable environment this is little enough risk and Single Attached Stations (SAS) are much cheaper than a DAS-oriented system unless the extra cost is justified on mission-critical grounds.

"The catch with deploying a system based on a synchronous access scheme is that the best way to get the best server-to-server performance out of it is to keep workstations off it."

The catch with deploying a system based on a synchronous access scheme is that the best way to get the best server-to-server performance out of it is to keep workstations off it. The reason for this is that workstations often generate large numbers of small frames, not to mention a significant proportion of broadcast traffic. This is precisely the sort of load which can ruin the performance of a synchronous access scheme by forcing each station on a token based system — which is how most synchronous access schemes work — to waste time processing large numbers of frame headers on frames, which are of little or no interest to most stations, and which flood the ring delaying the transmission of server-to-server traffic. Hopefully this server-to-server traffic is in chunks which are an integer multiple of the maximum data payload made available by the access scheme being used. It is often much more efficient to dual home the servers to both a synchronous access LAN and an asynchronous access LAN — say switched Ethernet — than to connect workstations to a server backbone via a bridge.

In short, network systems need to be tailored to the applications which they are intended to support if they are to be cost effective, efficient and ultimately reliable.

Graeme Le Roux is a Director of Moresdawn Pty Ltd (Bundanoon, NSW) and specialises in local area network consulting services.

Introducing LANtastic® 6.0. For the most complete office network solution.

Everyone in the office can run computers. The printers. The fax (okay, Bob's trying). But despite the wires, your office still doesn't connect. Before you throw up your hands, throw around this idea: new LANtastic 6.0.

The LANtastic 6.0 network operating system combines state-of-the-art networking with e-mail, faxing and scheduling. So you can communicate around the office (or around the world) right at your desk. No third-party software. No time-wasting steps.

LANtastic® is the first network connecting people. Not just computers.

What's more, a LANtastic system can operate peer-to-peer, as a dedicated server, or anywhere in between. (That's flexibility.) All at speeds quicker than LANtastic 5.0 for even faster printing and file sharing.

Now that's what I call productive, you say. But you already have an enterprise network? LANtastic 6.0 is the ideal workgroup solution. With our universal client, you can access the broadest range of networks—including servers from Novell, Microsoft, IBM and others.

And installation? Easy, too. LANtastic 6.0 is intuitive and reliable. So it's a breeze to learn and maintain. (Even for Bob.) And the price makes thoughts of working together even nicer.

So learn how to really connect your office. Call an Artisoft reseller or distributor.

Artisoft Australia 008 808432
Digital Solutions Pty. Ltd. 008 804635
Merisel Pty. Ltd. 008 251076

LANtastic by

ARTISOFT®

A better way to network people.

WAN VIEWS

Multimedia Over ATM? It's All a Matter of Timing

Steven Taylor

One of the cardinal rules governing the creative use of networking jargon is the principle of guilt by association. Two terms repeated together often enough become inextricably linked, regardless of the real basis for the connection.

One such forced pairing involves two of today's hottest technology properties: multimedia and ATM. There's no denying multimedia applications and asynchronous transfer mode transmission are going to enjoy a long and happy association in coming years. ATM's enormous potential as a bandwidth provider is a good match for multimedia's equally enormous potential as a bandwidth consumer. But run that concept through the jargon mill enough times, and 'good match' becomes 'only alternative.' And that just isn't the case.

Part of the problem is rooted in confusion over what exactly constitutes a multimedia application in the first place. Walk into your local PC store, and you'll find any number of 'multimedia upgrade kits' that bring some kind of audio-visual powers to desktop machines, usually in the form of CD-ROM drives and sound boards. But talk to a videoconferencing vendor, and multimedia takes on a completely different meaning, one that involves real-time exchanges of voice, video, images, and applications across the wide area.

Although all multimedia applications involve the delivery of two or more different information types (data, voice, and video), what really matters from a networking standpoint is the time sensitivity of a given application. In this regard, multimedia applications fall into two distinct categories: batch applications, which are not time sensitive on the millisecond scale, and interactive applications, which require timely delivery.

Like batch computing, in which a 'job' is submitted to the mainframe for computation and the results are returned anywhere from seconds to minutes to hours later, batch multimedia systems involve no real-time interaction among users. For instance, an insurance company might have a multimedia application that accommodates photographs or full-motion video displays of insured property, voice annotations describing the property, and spreadsheets or databases containing information about the property. These different data types are stored locally or at a central site for access by end-users. A delay of a few seconds is no cause for concern, as long as buffers are in place to take care of the biggest anticipated delay.

In contrast, true interactive multimedia involves real-time exchanges of two or more traffic types, as in videoconferencing. Because they are time sensitive, interactive applications require immediate access to sizable chunks of bandwidth. With the batch insurance application, users can wait until the entire files are delivered to their workstations; that kind of delay would destroy an interactive multimedia application like videoconferencing.

The basis for the assertion — let's call it the myth — that multimedia absolutely requires ATM comes from two factors: processing speed and a delay phenomenon called freeze-out. The speed requirement is simple. Because cells are easier to process than variable-length frames, cell switching is a must for high-speed transmissions. The exact point on the speed scale at which cell switching becomes essential depends on the processing power built into switches. A decade ago, T1 (1.544Mbps) transmissions required cell relay; today, the cutoff point is about OC-3 (155Mbps).

Freeze-out is a bit more complex. Since packet switching uses serial transmission (information is transmitted one bit at a time), a single packet occupies the entire transmission facility while that packet is being sent. The duration of this freeze-out is equal to the packet length (in bits) divided by the transmission speed involved.

Interactive multimedia is highly sensitive not only to transmission delays but also to variations in transmission delays. For instance, a 64Kbps pulse code modulation (PCM) voice channel — the long-time industry standard for voice traffic — assumes that an 8-bit byte of voice information will arrive once every 125 microseconds (or 8,000 times per second). If there is a significant variation in the delivery delay of a PCM byte, the algorithm starts looking for information that isn't there — hence the freeze-out.

The most likely source of this delivery delay is that the transmission link is occupied by a very long packet of data. The concern is that the transmission of very long data frames will cause unacceptable delays for delivering time-sensitive voice and video information. This leads to the blanket assertion that only ATM, with its short, fixed-length cells, can handle multimedia.

At this point, a couple of clarifications need to be made. First, freeze-out has an adverse effect on real-time applications only; batch multimedia has no problem coping with delay variations and in fact can gain in efficiency from the use of larger packet or frame sizes, since fewer overhead bytes are required. Second, freeze-out is a problem only at relatively low link speeds. For instance, the freeze-out time caused by a 4,000-byte (32,000-bit) data frame — which is approximately the maximum size allowed under frame relay — is about 0.5 of a second on a 64Kbps link (32,000/64,000 = 0.5).

"For non-interactive multimedia applications, such as imaging systems, frame-based technology may also be an attractive option."

That's clearly cause for concern. But at higher speeds, freeze-out isn't much of a problem. At T1, a 4,000-byte frame causes a freeze-out of just under 21 milliseconds. At T3 (45Mbps), freeze-out is less than a millisecond, and at OC-3 it's less than a quarter of a millisecond. Given these qualifications, then, freeze-out is a problem for interactive multimedia applications running over low-speed WAN links. At lower speeds, frame-based systems can combat the freeze-out problem by using smaller frame sizes.

Multimedia applications are evolving rapidly, and so is the need to support these applications across the wide area. ATM is certainly an option for handling these applications, but it is not the only option. For non-interactive multimedia applications, such as imaging systems, frame-based technology may also be an attractive option. And even though interactive multimedia is more readily supported by ATM, several frame relay-based interactive voice and video products are now available.

So the next time someone makes a blanket statement about multimedia and ATM, remember: that blanket probably isn't covering the whole truth.

Steven Taylor is President of Distributed Networking Associates, based in North Carolina.

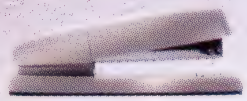

We're talking about your network. The thing that holds your business together.

But if you're like most companies, your network has gradually become a random patchwork of products and technologies, held together by a hope and a prayer.

And one day, it's going to unravel.

What you need is a strong, cohesive network infrastructure that will grow and change as your business does. You need the network fabric from SynOptics.®

The network fabric goes beyond our full line of intelligent hub solutions. It takes your hubs, routers, bridges, switches—even your cabling—and makes them work as one. Managing the entire fabric is our Optivity™ software, giving you a level of control and security you've never had before. The result is a network infrastructure nobody else can offer. Nobody.

But one company can't do it all. That's why we've teamed with other companies like Novell, IBM and Intel—allowing you to build your own seamless network fabric with the most advanced products available.

To learn how the network fabric can work for you, call 1 800 81 7070 in Australia or 0 800 44 7070 in New Zealand. Because when it comes to networks, no one's got it more together than SynOptics.

SynOptics

Building the Network Fabric

What's holding your network together?

INTRODUCING THE SIEMON CABLING SYSTEM

A comprehensive, quality-driven strategy for the physical layer. The Siemon Cabling System complies with the TIA/EIA, AS3080, & GOSIP standards and pushes beyond—offering unique product benefits and services such as our state-of-the-art cabling training facility. For a truly open cabling system that offers outstanding performance, quality, and value, specify Siemon.

Category 3, 4, 5 UTP Cabling Solutions
Fibre Connectivity
Structured Cabling Training Facility
Certification Program
Extended Warranty

FOR MORE INFORMATION CALL
PHONE (612) 905-9468
FAX (612) 905-9471

Quality • Service • Innovation • Value

The Siemon Company • 29/9 Powells Road • Brookvale, Sydney NSW 2100, Australia

Industrial Strength Integration

The GigaHUB provides a scalable and useable platform for the present while providing the framework for migration to new network technologies.

There are times when a network designer is faced with a requirement for a single hub to fulfill a number of roles like forming the centre of a collapsed backbone, combining a number of required features, and leaving room for expansion. These are apparently simple requirements but they often are difficult to meet at a realistic price. The result is a mixed vendor solution and/or the use of mixed technology and the result is nowhere near as neat or simple as most hub vendors would have you believe. Such mixed vendor solutions are usually costly to deploy and maintain. It is the reduction of such costs which has driven the push for integration around the world and led Fibronics to develop the GigaHUB as an adjunct to its successful MultiHUB range.

The GigaHUB is a deceptively utilitarian device, even when compared to Chipcom's excellent ONCore hub with which it competes, and certainly when compared to its other major competitors. But Fibronics' unit gives little if anything away in terms of performance. The GigaHUB is a twelve-slot chassis designed purely for rack mounting. A smaller 'Midi-GigaHUB' chassis is planned for later this year. The chassis also accommodates up to four 270W load sharing power supplies providing N+1 redundancy. The front of the unit incorporates several status LEDs, four push buttons for local setup, control and testing, and a clear four-line LCD display. The LCD display is also back-lit in recognition of Murphy's 42nd law of network maintenance; all equipment racks are placed such that there is insufficient light to read an LCD display.

GigaHUB is an adjunct to the MultiHUB range for sites which require higher performance and port densities than the older unit can provide and by the use of an adaptor cage (either 4- or 8-slot), it can accept MultiHUB adaptors. GigaHUB maybe configured with a number of native LAN adaptors for both Ethernet and Token Ring (Ethernet 10Base-T/2/5 with or without security, Token Ring UTP/STP and fibre) now and a variety of switching (Ethernet, Token Ring and FDDI), FDDI (F/O and TP SAS, FO DAS and a DAS to ATM card) and ATM (FDDI DAS to ATM, ATM/SONET at 155-Mbps) adaptors have been announced. There is also a range of channel adaptors available for those with IBM, Wang or other such environments. In fact the channel-to-LAN area is one in which Fibronics has been a leader for some years.

GigaHUB modules are hot swappable and can function with or without additional management modules. An interesting feature of the GigaHUB is that management of the unit is distributed among a range of modules which are physically daughter boards residing on cards and (therefore do not take up hub slots), and which may be installed in a variety of redundant configurations eliminating single points of failure.

Both in-band and out-band management is supported via Fibronics InterView SNMP Manager or any other generic SNMP manager supporting the relevant MIBs. This modular approach is carried to the GigaHUB's media interfaces which are provided via a number of 'Extension Modules' which are effectively transceiver modules. Up to four Extension Modules may be plugged into a 'LAN Concentrator' card which occupies one of the GigaHUB's twelve available slots.

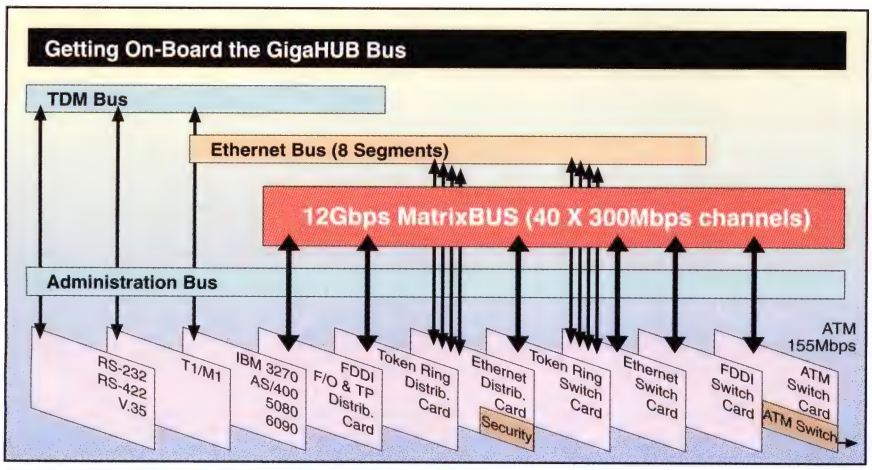

PRODUCT SUMMARY

Name: XH150 GigaHUB System

Description: High performance, ATM-ready, LAN distribution and switching hub with voice/data integration capability.

Price: XH150 GigaHUB Base system (inc. chassis and power supply) is priced at $12,748 and can be configured with almost any mix of about 30 modules.

Vendor: Fibronics International

Distributor: Olivetti Australia, 4th Floor, 140 William Street, Sydney NSW 2011 Tel: (02) 368 1988

As can be seen from the list of adaptors which so far have been announced for the GigaHUB, Fibronics clearly intend it to be their initial ATM platform and to make that possible the company has built the unit around a high performance 'MatrixBUS' architecture.

Catching the Right Bus

The MatrixBUS is a high performance switching backplane which provides 40 300-Mbps channels for use by switching LAN Concentrator Cards and ATM cards. The aggregate throughput of the MatrixBUS is 12Gbps which puts the GigaHUB in roughly the same class as ChipCom's ONCore, the SynOptics 5000 and the new Cabletron MMAC-PLUS.

There are actually three types of bus built into the GigaHUB: MatrixBUS, eight Ethernet busses (4 synchronous (ESBs) and 4 10Base-2 (EB2s)) and several service busses which are reserved for management connections, etc. The distinction between the GigaHUB's ESBs and EB2s is that the ESBs provide a single repeater count for the entire hub, while the EB2s are effectively four individual 10Base-2 segments which, when interconnected, count as one or more repeaters. This sort of configuration is not uncommon in the hub market, however Fibronics's clear distinction between the two is refreshing.

In summary, the GigaHUB should provide a reliable, scalable and useable platform for several years, which in turn provides network managers with the ability to lay the groundwork for migrating legacy systems without wondering whether or not their hubs will cope with such changes.

Graeme Le Roux

Making the Office Accessible

3Com's AccessBuilder provides network access to a large number of nomadic or small office users without compromising security.

Interconnecting several small network sites which are too small to justify dedicated communications lines or connecting such small sites to an enterprise WAN is often difficult to do cost effectively. Providing transparent network access to small numbers of remote users, let alone 'nomadic' users — those who might need to access a network from hotels, airport lounges and the like — can be a nightmare from both a technical and an economic point of view.

Modems and single user remote access software packages are cheap and simple to install, but can result in gaping holes in system security. On the other hand, products such as Microsoft's Remote Access Service are not necessarily a complete answer either.

If you have to cater for more than a few concurrent users a dedicated server is required and in cases where the software in question does not perform a bridging or routing function a large amount of unnecessary traffic is transmitted across the relatively low speed modem links. Other problems include security of such software and the fact that these solutions generally do not accommodate multiprotocol environments well.

The most common approach is to install all the required protocols on the server and have it translate traffic from the remote site into a protocol that the target host or server can understand — which is both clumsy and slow in many cases. 3Com's new AccessBuilder addresses all these issues.

AccessBuilder is the first fruit of 3Com's acquisition of Centrum Communications. It provides cost effective network access to a large number of nomadic, after hours or home-based users, or small offices requiring little more than e-mail access and intermittent file transfer capabilities without compromising security or creating an impossible support load. Best of all the solution is scalable, can be quickly deployed and could be used in conjunction with telco services such as 1800 and 13 numbers to centralise control of telecommunications costs.

Boundary Routing Similarities

From a functional point of view AccessBuilder has some architectural similarities to 3Com's Boundary Routing System Architecture in that it divides the basic functions of a brouter across two distinct devices: the AccessBuilder Remote Access Server and a software client which resides on the remote PC (or Mac). The client software is installed in a standard NDIS, ODI or PDS (Packet Driver Specification) environment with standard network client software (NetWare, LAN Manager, LAN Server, Windows for Workgroups, VINES, Pathworks, Chameleon, Wollongong Pathway, FTP's PC/TCP, WRQ's Reflections and other TCP/IP clients with SLIP or PPP support). Apple's ARA 1.0 and 2.0 are also supported over Ethernet and no modifications are required.

The AccessBuilder client acts exactly like a network adaptor driver, but drives a serial port and modem rather than a NIC, while the server itself is connected directly to the LAN, much like a conventional router or bridge.

Most popular modems are supported and others can be supported via user editable text scripts — this means all the features of your favourite modem, such as compression, can be used (the AccessBuilder has sophisticated header and body compression algorithms built in). The user sets a phone number and, when connected to the AccessBuilder, is required to provide a password before logging on to the network proper. Call-back, resource access control and PPP security (both PAP and CHAP) are all supported by the AccessBuilder and a security event log is provided along with an API set which allows integration with third party security products such as Kerberos servers. The AccessBuilder is managed via SNMP (MIB II with Centrum extensions), a console port or one of its asynchronous ports.

To filter unnecessary traffic from a link and to allow LAN-to-LAN access, the AccessBuilder supports IP and IPX routing and bridges AppleTalk, DECnet, XNS, NetBIOS, NetBEUI and other protocols.

The AccessBuilder is a simple device intended to provide only basic bridging and routing functions; if you want a sophisticated brouter with double overhead bells and whistles 3Com offers its NetBuilders. But while the AccessBuilder's brouting features are intended to be simple and basic they are in no sense incomplete. The unit supports IP, UDP, TFTP, ICMP, BootP, APR, RIP (IP and NetWare implementations), SLIP, PPP, IPX, SAP and over Ethernet only ARAP, DDP, NBP, RTMP, ZIP, AEP and AARP.

The AccessBuilder is designed to fit on a desk or in a single unit of a standard 19 inch rack. It is capable of supporting up to two interface modules in any configuration. Asynchronous modules support speeds of up to 115.2Kbps per port while synchronous ports (RS-232, RS-449, X.21, V.35) support speeds of up to 2.048Mbps per port. Network interfaces on the Ethernet model are AUI, 10Base-T and 10Base-2. The AccessBuilder for Token Ring supports both DB-9 and RJ-45 connection at either 4 or 16Mbps.

Graeme Le Roux

PRODUCT SUMMARY

Name: AccessBuilder

Description: Remote access server supporting sync and async connections to an Ethernet or Token Ring network by 'nomadic' users and small remote sites.

Price: 3C7513 AccessBuilder Base System (Ethernet) $3,743; 3C7515 AccessBuilder Base System (Token Ring) $5,243; 3C75100 AccessBuilder 16-Port System Bundle $9,443.

Vendor: 3Com ANZA, Level 7, 99 Walker Street, North Sydney, NSW 2060
Tel: (02) 959 3020

Lightening the Management Load

Cabletron's Spectrum 3.0 distributes the processing burden across multiple management servers.

Cabletron has it that the most effective way to oversee an oversized internetwork is to distribute control among multiple management stations. Hence its new Spectrum Advanced Management System 3.0 divides up management duties among a series of servers, each responsible for all devices in its domain. Polling takes place locally, rather than over WAN connections and all information is processed by the local server, rather than by one central-site console.

The key to Spectrum 3.0 is its distributed design. Spectrum software resides on multiple machines, called Spectroservers, that can be located across the corporate internetwork. Each server gathers configuration information about its domain and processes all management data on the spot and automatically backs up processed information. Spectrographs are client consoles that graphically display what's happening at the servers and from which users control their activities. Each viewing console can be teamed with one or more Spectroservers.

Servers can access information about devices under the purview of another Spectroserver by remotely capturing the other unit's screen via X Windows. Servers communicate via Cabletron APIs.

Spectrum software can operate under several versions of Unix and consoles aren't confined to the same operating system as the servers they're attached to. For instance, a Spectrograph running AIX can work with one Spectroserver running SunOS and another running Solaris.

The package's configuration capabilities go a step beyond what's now available in competing platforms. While some products can recognise a specific model of a vendor's router, for example, Spectroservers can pick different configurations of that router.

Spectrum's intelligent event-handling system now issues alarms only for root problems. When a router crashes, for example, only that failure is reported, rather than a message about every device affected by the malfunction. Spectrum incorporates the Remedy trouble-ticketing system from Remedy Corporation.

Getting Granular

Spectroservers do much of the system configuration on their own. Using an autodiscovery application, each server polls its domain to find the IP and MAC addresses of all attached hubs, routers, and other gear. The server creates its domain map of object-oriented data using SNMP MIB information collected from the equipment; Spectrum recognises all SNMP MIB II elements.

To extend Spectrum beyond the scope of SNMP MIB II, users can purchase product-specific add-on applications, called Management Modules, that supply MIB extensions for specific third-party devices. Management Modules are now available for 61 products and other add-on applications are available that add support for SNA and DECnet Phase IV.

While it's true that some other packages support MIB extensions for specific vendors' devices via third-party packages, not all extend the platform's entire range of services to those devices.

Like many other management platforms, Spectrum can link to key third-party databases. Using a two-step approach, managers can create reports about network activity. First, Spectrum management data can be exported to a third-party package such as the Ingres database from the Ask Group. From there, the information can be incorporated into spreadsheets or report writers. Another option is to send ASCII text from a Spectroserver to a spreadsheet or reporting tool.

Spectrum also works with major third-party management packages. Users can import information collected by OpenView from Hewlett-Packard, Sunnet Manager from Sunconnect, and IBM Netview/6000.

Cabletron says it hopes by year's end to launch management applications in conjunction with a rollout of new hardware modules for its hubs. Also on the drawing board is a version of Spectrum for Windows NT.

Mary Jander

PRODUCT SUMMARY

Name: Spectrum Advanced Management System Version 3.0

Description: A management system for corporate networks based on distributed architecture, spreading control of devices across multiple servers and consoles.

Price: Spectroservers $15,000; Spectrograph $8,000; Management Modules from $1,800.

Vendor: Cabletron Systems, Unit 8, 25 Frenchs Forest Road, Frenchs Forest, NSW 2086 Tel: (02) 950 5900

Spread the Views

Cabletron's Spectrum 3.0 system distributes net management chores across multiple Spectroservers, which oversee all devices in their domains. Spectrographs display the activities of one or more servers using X Windows; consoles needn't run the same operating systems as the servers they're teamed with.

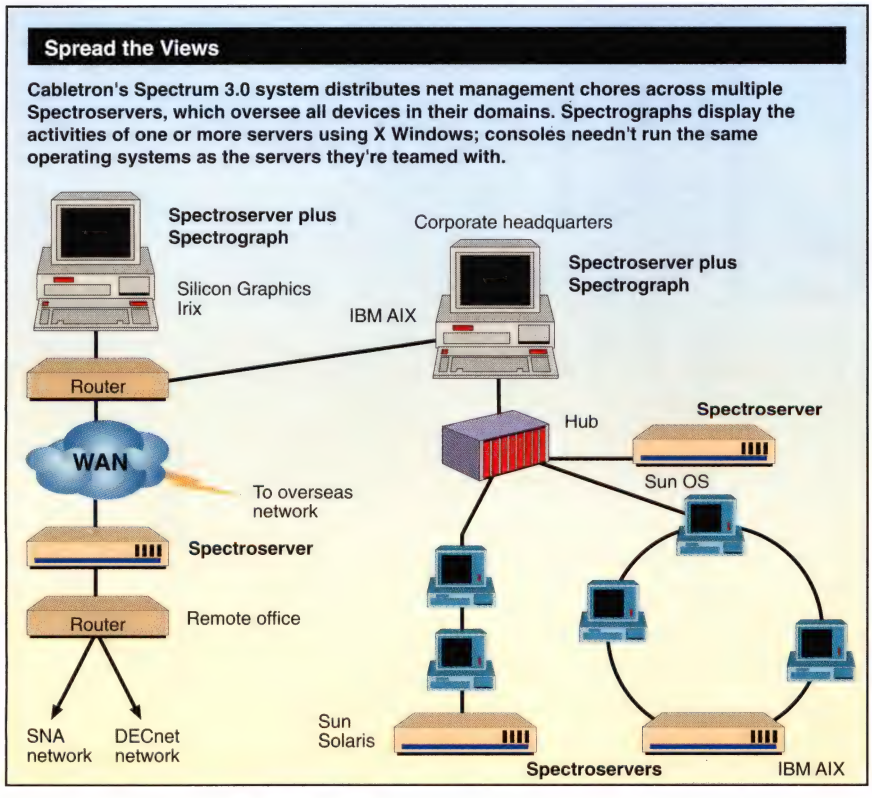

Tony Pool and Malcolm Hamilton
invite you to

MobileCom
EXHIBITION

Australia's 1st Mobile Communications Event

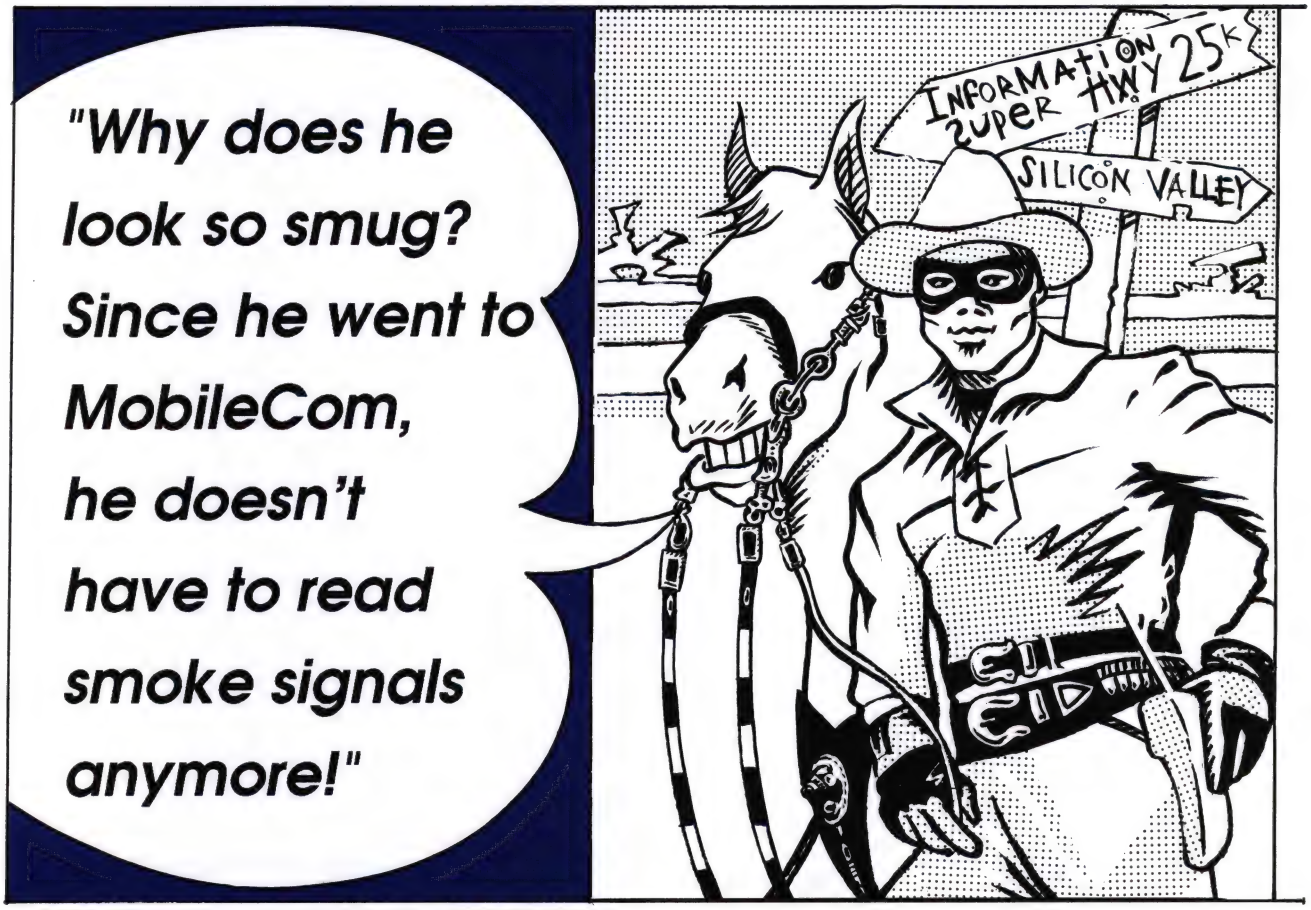

OPEN TO EVERYONE – TRADE AND PUBLIC

SYDNEY	Darling Harbour Friday 25th-Sunday 27th November 1994
MELBOURNE	Royal Exhibition Building Friday 31st March-Sunday 2nd April 1995
BRISBANE	Brisbane Convention and Exhibition Centre Friday 16th-Sunday 18th June 1995

HIGH PROFILE EXHIBITIONS PTY. LTD.

High Profile Exhibitions Pty. Ltd. A.C.N. 063 398 836 1st Floor, 21 Swan Street, Richmond 3121 Victoria Australia
Phone (03) 428 0415 · (03) 428 0417 · Fax (03) 428 9508 · Mobile (015) 356 698

When Two's No Longer Company: The 1997 Review

The Government's recently announced carrier review could result in sweeping changes to the *Telecommunications Act*. Peter Waters considers the issues which need to be addressed.

A review of the *Telecommunications Act* of the scope announced by Communications Minister Lee was not required to farewell the wireline duopoly. Section 57 of the Act already entitles the Minister in his or her 'absolute discretion' to grant carrier licences, without any restriction on the number. The Government could simply wait until expiry in 1997 of its contractual restriction with Optus (and Vodafone in relation to mobile carriers) on the issuing of more licences.

Clearly, the Government has taken the opportunity presented by the end of the duopoly to completely review the current industry structure. The result could be a *Telecommunications Act*, and quite possibly a *Radiocommunications Act* and *Broadcasting Services Act*, which look very different from today's legislation.

Why the Duopoly?

While we are to bid farewell to the fixed network duopoly (some with more fondness than others), it is still useful to review the duopoly's policy basis when considering what should replace it. The Government saw the limited duopoly period as providing the new wireline carrier with a strong incentive to rapidly roll-out an extensive network to maximise the advantage it would have over subsequent entrants on expiry of the duopoly. Telecom also would get the message that it had a limited opportunity to 'shape up' before the full force of competition hit.

A staged transition to competition was considered more likely to successfully break Telstra's entrenched position than an immediate switch from monopoly to full competition. A single, strong player would be a better match for Telstra than a 'gaggle of pygmy carriers' which Telstra could play off against each other. Giving preference to one new entrant also may limit uneconomic duplication of infrastructure as two fairly large networks would already be in place when subsequent new entrants came to consider their investment plans.

There is likely to be debate during the course of the review about whether this staged competition approach actually achieved these goals. Telstra, no doubt, would take the view that the 'most tilted playing field in the world,' as Frank Blount terms it, has succeeded in producing a vigorously competitive environment in which the last vestiges of its long held monopoly have just about been expunged.

Optus, on the other hand, would probably assert that the limited share of the market Optus has been able to achieve is the result of the regulatory safeguards. Unlike with the initial debate about the duopoly, the position which each carrier takes on the post-duopoly world will not be determined solely by their assessments of their relative positions to each other. They will also need to consider the impact of potential new entrants — including their 'shared' circumstance as the existing carriers.

The key issues to be considered in the review are:

- Whether the rights and obligations which attach to carrier licences should be changed;
- Whether there should be different categories of carrier licences, and whether there should be limits on the number of licences in any category;
- The future of the Universal Service Obligation in a new multi-carrier environment;
- How interconnection should work, given that potentially it will be more complex than the current bi-lateral arrangements between each of Optus, Vodafone and Telstra; and
- What criteria should be imposed on new carriers, including any restrictions on foreign ownership levels.

Some directions already seem clear. There have been press reports that a draft paper prepared by the Minister for presentation to the ALP Conference in Hobart commits the Government to high levels of Australian ownership in the industry. As Vodafone is largely foreign owned and Optus has a significant level of foreign ownership, there may not be much more room for foreign owned carriers.

The Minister also stated that the review would take into account the Hilmer Report, which recommended against specialist regulators in favour of a single competition regulator — basically a 'super' Trade Practices Commission.

Two issues are likely to be a particular focus of the review: first, whether the current land use powers and environmental exemptions are politically sustainable, if there are many more carriers; and second, what will the defining differences between carriers and service providers or resellers.

How Green?

When the 1991 *Telecommunications Act* was enacted, there was a concern about extending to a privately owned carrier the powers which Telecom traditionally has had to enter and locate facilities on land. However, the bidders successfully argued that without similar powers they would never be able fully to exploit the window of opportunity provided by the duopoly, and that Telecom would unfairly benefit from the years of monopoly when it used those powers to build its network. As with Telecom in the past, the new general and mobile carriers also were granted exemptions from State environmental laws, subject to their compliance with a special Federal code.

While vesting powers of this nature in just one more general and one more mobile carrier in addition to Telecom was politically saleable, the prospect of yet more parties having these powers is much more troublesome. The political difficulties have been exacerbated by the increasing public disquiet over mobile towers.

Alive to these concerns, the Minister has stated already that hordes of carriers will not be descending to dig up our front yards. This suggests that the current powers may either be heavily restricted or eliminated altogether, with carriers having to negotiate with landowners to obtain easements etc. This raises a question of whether only new carriers would be subject to these reduced powers, or whether the existing carriers also would lose their current powers. Optus and Telstra have put in place significant networks using their current powers, and loss of those powers, at least in relation to the network facilities built and planned on their assumption, could cause substantial interruption and dislocation to their operations. However, if only the existing carriers retain these powers, the new carriers are likely to argue that they are being treated as 'second class' carriers.

In the United Kingdom, the Government decided that new carriers would not automatically be given statutory land access powers, but an assessment would be made for each licence. The Government saw a relationship between these rights and the obligations and commitments which the new carrier was assuming; principally the amount of network construction the new licensee would undertake and the extent of service coverage it was proposing. In other words, did the new entrant really 'look like a real carrier,' or was it really a reseller with no or a narrow facilities base?

Carrier vs Service Provider

Determining what should be the regulatory differences, both rights and obligations, between a carrier and other suppliers of telecommunications services presents the industry with an even thornier set of issues.

Why keep the current regulatory distinction between carriers and service providers at all? Even in a market with many more carriers, there probably still would be a range of suppliers of telecommunications services who would not want to be treated as carriers. While service providers are subject to a fairly light-handed level of regulation, carriers are subject to an array of obligations: carriers are required to file tariffs, they are subject to rules limiting differential pricing between customers, they must institute and comply with customer complaint mechanisms, they must participate in and fund the Telecommunications Industry Ombudsman scheme and, most importantly, they must make USO contributions.

While service providers supplying 'PSTN-like' basic voice services probably would regard the benefits of carrier-like status as outweighing the burdens, many service providers would not, particularly those providing specialist services or serving narrower customer bases.

Demarcation Disputes?

The major complaint of some service providers with the current regulatory regime is that carriers are only required to treat them the same as any other customer, while carriers are required to treat each other preferentially. In particular, carriers and service providers currently have access to very different levels of functionality in the services they take from other carriers. Through interconnection rights, carriers can seek service and facilities from other carriers which the supplying carrier would not otherwise supply outside its own operations.

By contrast, service providers are only entitled to the tariffed services. Each carrier determines the nature and configuration of the basic carriage services which it puts on the market, and the terms on which the service is provided, including technical standards. The effect of Austel's BCS Opinion was that carriers did not have to separate out and, therefore, make available to service providers, advanced network functions like call forward and call back.

The policy basis of this vertical functional separation between carriers and resellers was set out in Austel's 1990 Resale Report:

'The second carrier's task of establishing a comprehensive alternative national network will require substantial resources — both capital and human resource . . . Permitting resellers to directly compete with carriers could limit significantly the second carrier's ability to generate enough surplus on some parts of its operation to fund such commitments.'

A fundamental question for the review will be whether this 'functional distinction' between carriers and service providers should be adjusted, and whether the service providers' rights to require carriers to supply functionality should be expanded. This question cannot be addressed solely on the basis of what level of 'unbundling' of carrier networks is technically feasible, but also on major economic and policy issues. As the current unbundling provision of the Act identifies, unbundling requires consideration of the impact on the service provider's ability to develop and offer products and on the carriers' ability to exploit fairly the economies of scale and scope which they derive from the investment they make in network infrastructure.

The carriers may assert that investment incentives will be undermined if carriers do not have an assurance that they can earn a reasonable return on their investment, or if they are required to make available to their competitors any functional advantages they gain from investment in new facilities. The end of the duopoly does not necessarily mean the end of the need to provide incentives for network construction amongst the facilities-based carriers, however many licensees there may be after 1997.

Service providers are likely to respond that 'open networks' are the way of the future, particularly with the potential explosion of information and entertainment services, and that the carriers derive enough advantage from their ownership of the infrastructure without the protection of functional bundling.

An apparently easy way of getting around the problem would be to expand the category of 'carrier' to include 'PSTN-like' service providers who are currently operating under the 'service providers' banner. However, this course would only transfer to the carrier interconnect arrangements the problem of trying to provide incentives and rewards for facilities investment by some carriers, while at the same time ensuring that other carriers which have chosen to have limited facilities of their own get the access they need to required services on a fair basis.

Telecommunications Acts have a habit in Australia of turning over faster than the underlying technology. The 'straws in the wind' suggest that we could have a radically different industry structure after 1997.

Peter Waters is a partner at Gilbert & Tobin, and specialises in technology and telecommunications. This column sets out his personal views and not those of the firm or its clients.

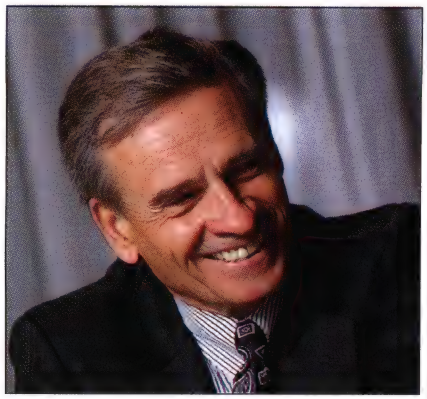

TCI, Convergence and Cable Competition

Fred A. Vierra
Executive Vice President
Tele-Communications, Inc.

As one of four members of the Chief Executive Office for Tele-Communications Inc. (TCI), Fred Vierra is responsible for the company's global portfolio which includes cable TV systems and programming. He serves as a TCI board member for Australis Media, TeleWest, Turner Broadcasting Co. and The Discovery Channel among others. Prior to joining TCI in 1991, Vierra was President and Chief Operating Officer for United Artists Entertainment Co. (UAE) and President of United Cable Television Corp., which was merged into UAE in 1989. While in Sydney to address the Pacific Telecommunications Council (PTC) mid-year conference last month he spoke with Liz Fell.

In your speech at the PTC conference, you described TCI's mission as bringing 'telecommunications into tomorrow.' But isn't it correct to say TCI is still predominantly a cable TV company in the US?
Vierra: Yes, if you don't count Teleport. Because of vertical and horizontal limitations, the Congress has more or less said to TCI and the other major cable companies: 'You're about as big as we're going to let you get.' TCI has had a compounded growth rate of 20% per year for two decades and it's not going to give up growth. So it has to grow in new revenue streams and new opportunities. On an international basis, I can continue to grow the business at the 20% rate, subject to availability of capital. I can literally sit on the sidelines and find out what works for Time Warner or what works for TCI and other cable companies and then export it to Europe, Asia and Latin America.

AC: *And to Australia?*
Vierra: Absolutely.

AC: *Can you describe the Teleport venture?*
Vierra: Teleport is an alternative access telephone company operating in some major cities. There are other cable companies in there: Time Warner, Cox, Continental and Comcast. It is a good strategic fit for TCI because it uses the fibre that the cable companies are building in these cities. We can put in the fibre that we're going to put in anyway and provide enough capacity for Teleport so that it can pick up the alternative access business. It has been very successful.

AC: *Why did the TCI-Bell Atlantic mega-merger collapse?*
Vierra: For a variety of reasons, not the least of which was the continuing rate re-regulation on the part of the Congress.

AC: *There were reports of differences between TCI's John Malone and Bell Atlantic's Ray Smith. Were these correct?*
Vierra: John Malone very, very much wanted to do the Bell Atlantic transaction. I can't speak for Ray Smith because I don't know him. Certainly there were a number of people in TCI who were not nearly as enthusiastic about this as John — myself included.

AC: *Why weren't you enthusiastic?*
Vierra: Personally, I thought there would be serious culture clashes. I work with telephone companies every day. We can make decisions at TCI very quickly. Telephone companies tend to be a little more deliberative in the decision-making process and that can be a little frustrating. The question was which culture would survive? Would we become bureaucratic and have to go through 42 iterations of analysis and be reviewed by several committees?

AC: *Bell Atlantic now appears to be moving swiftly to enter the cable TV market where it will compete with TCI.*
Vierra: My expectation is that within some reasonable period of time, and don't ask me to define it, TCI will probably compete with all of the Regional Bell Operating Companies [RBOCs] and maybe have various joint ventures with all of them. Let me lay a foundation for this. The Congress is considering legislation that would open the local loop, the RBOC turf, to competition. That may happen in this Congress or it may happen in the next, but there's little doubt in my mind that it's going to happen. At the same time, under certain circumstances and rules, they're going to let the RBOCs into long distance and into video.

So here we are with 10 million plus cable subscribers in the US, plus investments in another five million. We've got cable subscribers in every RBOC territory. If we go into the telephone business, and they go into video, we'll be competing in each market with each of them.

AC: *Does this mean there will be two pipes to every home or will there be interconnection to avoid duplication of infrastructure?*
Vierra: That's a good question. The answer is in the hands of the gods or if not in the hands of the gods, in the hands of the government. TCI's position has for a long time been that we ought to have one plant that the telephone company and the cable company share and compete head-to-head, albeit for the customers, using the same infrastructure. But I don't think that's what the Congress wants.

AC: *What does Congress want?*
Vierra: I'm sticking my neck out here and guessing, but I think they want the two pipes.

AC: *Given the massive investment needed to build this so-called superhighway to homes, isn't two pipes somewhat extravagant?*
Vierra: The only example that exists today is the United Kingdom where there are two pipes.

AC: *And TCI is building one of those pipes with US West.*
Vierra: Yes. We're in the United Kingdom, in Norway, in Sweden and in Hungary. We offer telephone and cable in the United Kingdom and cable everywhere else.

AC: *With the UK TeleWest venture, what proportion of revenues comes from telephony and from cable?*

Vierra: The revenue from telephone is about the same as the revenue from cable. I don't have the figures with me, but it's about 50:50. That includes business telephony as well.

AC: *Are you signing up many telephone customers?*
Vierra: Well, these numbers will be directionally correct though they are not necessarily correct for any one system. Roughly we get 30% cable penetration in new areas with the 'first pass.' That means the first time we knock on the door, 30% of people take cable and 20% take telephone. That's one out of five taking telephone, a big number. The revenues are something more than £40 per month per home when both are taken. That's about $US60 a month. Serious money.

AC: *I've seen a BT claim that the domestic phone line is used only for about five minutes a day.*
Vierra: Well, I wouldn't challenge BT's numbers. But we've always looked at it differently. Often the cable industry in the US is compared to the telephone industry for service and reliability. Yet the average television set is on about 7.5 hours a day and the average telephone is not used for a fraction of that. But it's really kind of academic. Have you ever looked at BT's revenues?

AC: *Not recently. But it must have lost share in the lucrative long distance market.*
Vierra: That's another element. Most of the telephone companies around the world have used long distance and international to subsidise the local. As you bring in a Mercury to compete with BT, those numbers change. You have to begin to raise the cost of the local service.

AC: *BT is trialling ADSL over its existing copper wires. Do you see that as a technology that has arrived too late?*
Vierra: I have to be careful what I say here. I think ADSL is a terribly limited technology. The twisted pairs in and of themselves are very limited and the telephone companies, including BT, are creating a lot of confusion in the market by talking about things they're going to do that are a lot more difficult to do than they're letting the financial markets know.

AC: *But ADSL is still at the trial stage.*
Vierra: They're all trials. We've got trials too. We have an experiment in the Denver suburbs called VCTV [Viewer Controlled Cable Television]. Our partners are AT&T and US West. There are two parts. There's near video-on-demand, 24 channels of pay-per-view scheduled so that the major movies are only 30 minutes away from start time, and then there is true video-on-demand.

AC: *Is AT&T providing the switching?*
Vierra: I asked Bell Labs how they were going to do that and they said: 'That will be bio-mechanical.' That means we've got guys in the room slapping tapes on tape decks! For a marketing test, you can do that. The question that is being asked here is as follows: if near video-on-demand gives you the same or equivalent buyer rates as true video-on-demand, then maybe you don't have to invest in the technology of file servers. But if movies-on-demand are a substantially higher revenue stream, then clearly you want to help develop the file servers.

AC: *Have you found a difference between movies-on-demand and waiting for them to start?*
Vierra: We have confidentiality agreements with AT&T and with US West so I have to be careful. But let me just give you one quick one. If you want to watch true movies-on-demand, it's a dollar more than near movies-on-demand. And we have found that many of our customers will pay the dollar not to wait 10 minutes for a start.

AC: *That's interesting. Rather than try a merger, US West has what is referred to as a 'partnership' with Time Warner. Is this the way forward?*
Vierra: US West has invested about $US2.5 billion in Time Warner Entertainment.

AC: *That's on the program and cable side of the business.*
Vierra: Yes. I think that partnership is a fair choice of words there and I think from US West's point of view, it was a very good decision.

AC: *There have been reports that TCI is having discussions with Viacom about acquiring their cable properties. Is this part of the fallout from Viacom's merger deal with Paramount?*
Vierra: Sumner Redstone [Viacom] has a big load of debt and he is trying to decide what makes sense to jettison. I'm not in those discussions but I will go on the record and say that we would be interested in those properties if at the end of the day Sumner decided to sell them.

AC: *How would Congress respond to TCI becoming larger, given it's already the largest US cable TV company?*
Vierra: In this whole process of re-regulation, Congress mandated to the FCC to establish horizontal and vertical limitations. We could do that, and maybe another one, before we started bumping up against the limits.

AC: *Why does one hear TCI referred to as a thug and as a strong-arm bully. Didn't Al Gore once call it the ringleader in the 'cable Cosa Nostra'?*
Vierra: Yes. And he referred to John Malone as Darth Vader! But let me just set the record

clear here. It's true that we have been accused of those things. But when you're not busy, check the source and see how many of them came from telephone companies who were trying to change the regulatory landscape and how many of them came from broadcasters who were concerned about their erosion of market share.

AC: *I assume TCI has been just as heavy in its own lobbying efforts in Washington?*
Vierra: Actually, I wish that had been true.

AC: *Why wasn't it? Was it a question of resources?*
Vierra: We didn't have the resources and I think we were terribly naive about the process. We're smarter now. That's because we've got arrows on our arse.

AC: *But when you look at TCI's history, these accusations appear again and again, even from local governments. Why is TCI seen as a bully? Or is it a bully?*
Vierra: It's a fair question. If you look at Televisa in Mexico, they're seen as a bully. Rupert Murdoch is often referred to as a bully. TCI is referred to as a bully. I've heard Time Warner referred to as a bully. I think any time you get a certain size, that is just the risk you run.

AC: *And TCI is now the world's biggest cable company. In the early 1970s, when TCI founder, Bob Magness, and later John Malone, started the company, why was it called Tele-Communications, Inc. Did the name indicate some ability to see the future?*
Vierra: I don't know the answer to that. As the business went forward, I think John Malone, in particular, who was very much a visionary, began to see the need to integrate vertically into programming investments and you can look at the pattern of investments that followed that.

AC: *And that pattern of investments led to charges of dominance which saw TCI spin off some program and cable assets into a company called Liberty Media?*
Vierra: Yes. Liberty is coming back into the company.

AC: *Some commentators would argue that Liberty never really moved out, that it is really a TCI controlled venture.*
Vierra: Tacky! Tacky! It had other shareholders.

AC: *But aren't most Liberty directors associates of TCI?*
Vierra: I don't think that's a fair statement. If you said the voting stock was controlled by Bob Magness and John Malone, then that's a fair statement. The directors were not predominantly TCI. But let me discuss the reorganisation of TCI. I'm not sure how much you know about that.

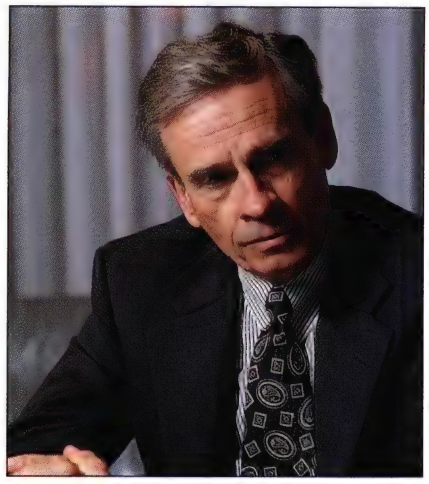

AC: *Very little other than that Liberty will be taken back into the fold.*
Vierra: Yes. And subject to that transaction closing, which is scheduled for the first week in August, we're going to reorganise the company into four subsidiaries or divisions: a programming entity, an international entity, a domestic entity and a technology venture capital entity.

AC: *What is the role of the technology venture capital entity?*
Vierra: In our business over time we see all technologies. Everybody comes to us. They go to a lot of other people too. Every now and then we see a technology that looks very attractive, so we've decided to give organisational recognition to the fact that we're making these investments by creating an entity that will be funded separately.

AC: *So would it be true to say TCI has an interest in most so-called convergence areas?*
Vierra: We have investments or joint ventures in all areas.

"It's true that we have been accused of those things [bullying tactics]. But when you're not busy, check the source and see how many of them came from telephone companies who were trying to change the regulatory landscape and how many of them came from broadcasters who were concerned about their erosion of market share."

AC: *What about the computing side?*
Vierra: We actually have made a venture capital type of investment in a computer technology company. But that's confidential and I can't say any more than that.

AC: *Not even the name?*
Vierra: No. I can't tell you that.

AC: *What about the converters and set-top boxes?*
Vierra: The investment we made would be in the computer chip for the set-top. We don't actually have an investment in a company that makes converters. But we own 49% of TSX, a hardware company that makes amplifiers and so on which is a very successful competitor to General Instruments in some areas.

AC: *Doesn't TCI have an agreement with General Instruments for the converters?*
Vierra: We gave an initial order for decompression converters and as part and parcel of that order General Instruments has to licence the technology to other manufacturers. The reason for the order was to give them the vehicle to raise the necessary capital to go forward if they needed that.

AC: *So you will not let this be proprietary in any way?*
Vierra: That's correct. We will not let GI have the technology on a proprietary basis. But we have more than one supplier. We have agreements with three different companies: Hewlett-Packard, Scientific Atlanta and General Instruments.

AC: *Were you involved in the Australis Media decision to choose General Instruments in Australia?*
Vierra: No. They made that decision on their own.

AC: *How does satellite fit into TCI's US operations?*
Vierra: We have Netlink which is a satellite-delivered service to C-band dishes and, in partnership with other of our colleagues in the cable business, we also have a venture called Primestar. Currently Primestar is mid-power but it will go to high power Ku-band in the very near future.

AC: *Will those be high power Ku-band DBS services similar to the Hughes DirectTv services?*
Vierra: Exactly.

AC: *Moving to Asia, I've read that TCI has entered Japan through a joint venture with Sumitomo. Will this involve cable and telephony?*
Vierra: The actual investment we have today with them is in CSN, Cable Soft Network, which is a general entertainment programming service. But the venture that you read about recently was to expand that relationship to do telephone and cable, subject to decisions of the Government.

AC: *Doesn't Sumitomo have cable operations already?*
Vierra: They've got a lot of cable franchises.

AC: *Have you invested elsewhere in Asia, leaving aside Australia and New Zealand?*
Vierra: So we're not counting Australia and New Zealand as Asia?

AC: *I was looking further north first! Are you involved in the race to get programming up on the Asian satellites?*
Vierra: The only place where we're directly involved is in Asian Business News. And I'm on the board of Discovery who went on Palapa and have committed to a transponder on Apstar. But that was a decision they made themselves.

AC: *What about on the ground in other parts of Asia?*
Vierra: On the ground, there isn't anything that's imminent. We are having discussions with a significant partner in India that's very attractive to us, but it would be premature to say that we're doing anything there. We'd like to.

AC: *Turning to Australia, why did you decide not to join the PMT [Packer-Murdoch-Telstra] Pay TV group?*
Vierra: In the early discussions with what became PMT, there were three of us involved: Time Warner, TCI and Comcast. At that time, under the rules, you could only have 35% ownership and no one company could have more than 20%. So when they divided up the equity, and we divided up the equity among the three US partners, it wasn't big enough to fool with. It had nothing to do with the partners and nothing to do with the venture. It was simply that the piece they were offering was too skinny.

AC: *Who were you talking with at that early stage?*
Vierra: Lynton Taylor and Kerry Packer from Nine and Rupert Murdoch. Telstra wasn't in at that point.

AC: *What is TCI's relationship with Murdoch?*
Vierra: Friendly. We're 50:50 partners in Request TV, a pay-per-view venture and, on a personal note, I hold him in the highest regard.

AC: *What is TCI's financial interest in Australis Media?*
Vierra: We own half of Gerry Lenfest's investment and we put in another $6.2 million. It's about 26.5%.

AC: *When I talked with you in Hong Kong last year . . .*
Vierra: Was I consistent with what I'm saying now?

AC: *Well, you said Lenfest was alone in the Australis Pay TV venture and that TCI was not directly involved. Further, you said you had been talking with the PMT group. Would you like to explain your change of mind since then?*
Vierra: Everything I said to you in Hong Kong was true and is not true today! An honest answer is that I frankly wasn't very enthusiastic about Lenfest's investment in Australis.

AC: *So did Lenfest talk with you before committing funds?*
Vierra: No. Gerry Lenfest made that decision absolutely unilaterally. But you're quite right. I was talking with PMT and I was also talking with Lynton Taylor [now with Cable Television Services] about a variation on the PMT theme. In the meantime, Gerry Lenfest and Rod Price from Australis came to Denver and met with John Malone and myself. When we really began to review the Australis venture and thought about it as being similar in opportunity to BSkyB in the UK, it began to look pretty attractive.

"Let me leave you with one final quote: 'Multimedia is the fastest growing zero billion dollar business in the world'."

AC: *But the UK market is much larger?*
Vierra: Well, today Murdoch has got between two and three million customers on the satellite. Australia has six million TV households, so it wouldn't be unrealistic to at least have programming in two million households here. That would be a very attractive business.

AC: *Aren't you bothered about the large upfront payment Australis/Lenfest paid for the satellite pay TV licence?*
Vierra: Initially that bothered me a great deal. But as I went through the business plan and did a back-of-the-envelope calculation, if you think about it as vertically integrated with the programming businesses too, it begins to look pretty attractive. The other thing that swayed me was that John Malone thought it was a good idea!

AC: *With Telstra now in PMT, do you see it as a formidable competitor?*
Vierra: From the outside looking in, PMT doesn't seem like a marriage made in heaven!

AC: *How does Australis, with satellite and MDS licences, provide a vehicle for TCI to offer video, voice and data?*
Vierra: Today I would say that our efforts are to develop programming and to develop a distribution system for that programming. As the telecommunications part of the business evolves, we'll watch it.

AC: *The telecoms duopoly disappears in 1997.*
Vierra: But what are the rules? And are they going to hold the existing telephone companies out of video? I just don't want to get into a war with the guy that's got all the guns.

AC: *Where is the real money going to be made in the future? In the software? In the distribution? As the gatekeeper?*
Vierra: In the US, we've been accused of being a gatekeeper because we have the subscribers and the case could be made that if you want to get access to the subscribers you have to deal with TCI and Time Warner. But look at the UK experience. It's such an interesting contrast because direct-to-home satellite was out there first and we're coming along building the cables second. Every distribution system has a channel captain controlling the channels. BSkyB is the channel captain right now in the UK, but that's changing.

AC: *In the end, surely TCI's interest in telecoms and TV would lead to a preference for the ground and cable systems?*
Vierra: I don't think it's that simple. There are so many things that come into play. If I'm uncertain about the stability of government, I'd rather be on the programming side because I still own the rights, though I may have the costs of a transponder. But compare that with the $US700 million that TCI and US West have invested in the ground in the UK. I wouldn't want to invest that kind of money, even with a partner, in an environment that I found politically unstable.

AC: *Does it matter that Japan, for instance, has just acquired a socialist prime minister?*
Vierra: I would make the same claim about our president!

AC: *I assume that these are exciting times for TCI, but they must also be rather confusing.*
Vierra: Do I find it a little confusing? Of course. I sometimes take a pretty myopic view of things. My attitude for the most part is that I know the basic architecture we're building is right. Where we are building plant, it's fibre to the node. Having done that, I don't really have to invest in all these multimedia.

AC: *Do you think the node is as far as fibre will go?*
Vierra: That's all it needs to do because coaxial cable has the same capacity as fibre if you don't have to send the signal too far.

AC: *So you are not jumping on the multimedia bandwagon?*
Vierra: Let me leave you with one final quote: 'Multimedia is the fastest growing zero billion dollar business in the world.'

Liz Fell is a freelance journalist based in Paddington (NSW).

WE JUST PUT MORE DISTANCE BETWEEN RLN AND OTHER REMOTE ACCESS SOLUTIONS.

For some time, Remote LAN Node has been acknowledged as the best remote node solution you can buy, letting users dial into their LANs and work just as if they were locally connected. And it's been proven in use across the world. Now RLN 2.0 has really left the others in the dust.

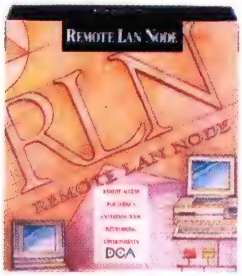

RLN comes preconfigured on a rack-mountable 486 server for 4, 8 or 16 simultaneous users, or as software-only for 2-16 simultaneous users.

RLN now supports Token Ring as well as Ethernet. In fact, you can access either one with the same client software. RLN is the only remote access solution that works in virtually any single or interconnected network environment.

There's also new SNMP support, and a built-in Windows™-based Remote Manager console for server management from any remote PC. You can also simultaneously update user information on multiple servers from a single location. And RLN has a Custom Client Build utility, so configuring your remote clients is easier.

yes It runs with NetWare

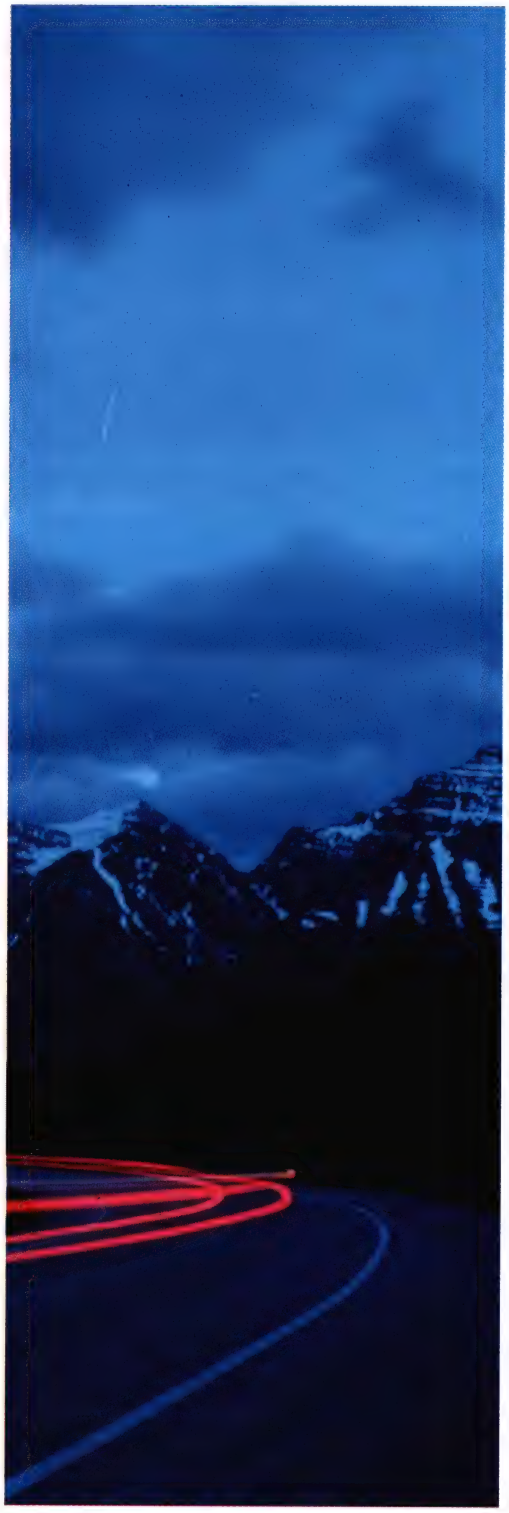

RLN has always included powerful security features. But now that we've added enhanced data encryption, password aging and support for Challenge Handshake Authentication Protocol (CHAP), it's the most secure solution you can buy.

So call for more information and an informative white paper on remote computing.

Like RLN users the world over, you'll see what lengths we've gone to, to make sure no other remote solution comes close.

02 451 8255

DCA

FAX TO DCA ON (02) 451 4034 TODAY

PLEASE TICK ☐ ***Yes!*** Please send information for remote Lan Node

PLEASE TICK ☐ ***Yes!*** I would like to upgrade to remote Lan Node

NAME: _____

TITLE: _____

COMPANY: _____

TEL: _____

ADDRESS: _____

POSTCODE: _____

Tom Amos

Franchising — A Tilt at the Future

If you are granted a special privilege in exchange for providing community goods, can you on-sell that privilege for a profit? This is a burning question that may soon be put to the test if the carrier plans to introduce exclusive franchise resellers to compete in the resale market actually get off the ground.

Ice-cream, cosmetics, hamburgers, furniture and mufflers all have a common thread; they interface to the public in a homogeneous way via organised channel distribution systems or privileged franchises. The public have come to accept over many years franchising of services as a way of determining quality and predictability of the service. If executed correctly everyone wins — or so the theory goes. But in the wonderful interventionist world of telecommunications where we have USOs and non-cost-based services, such simple theory may be just that — theory.

For some time now the three licensed service suppliers have been developing and operating marketing channel distribution strategies with varying degrees of success, particularly in the PMTS area, where there has been spectacular growth. Each supplier strategy has been different when viewed from the user perspective. But in the normal world all is set to change with the recent abolition of Strategic Partnerships (SPAs) and the introduction of more focused vertical dominant carrier tariffs. A change from the Ice Age to a period of enlightenment?

The use of resellers and the ability to discriminate between dominant and non-dominant suppliers has resulted in a number of competitive 'me too' tariffs, but this is set to change as all carriers move down the competitive market path to vertically organised channels with relationships that vary from intense to exclusive.

Now there are always ways around the intent of limited market power abuse that was encapsulated in the Act, although it took years for the lights to come on. Of course, declaring everything basic was a major set back.

There are many types of franchise, but the one that may be most suited to the telecommunications industry is the partial ownership exclusive vertical one. The products do lend themselves to repackaging, and, if cleverly done, it would be almost impossible to unravel who makes what and where. After all, there are already a dozen or so names for a local access line, so what's a few more!

Is there a benefit in privileged arrangements or, more importantly, do users get a better service at a lower price from the introduction of another mouth in the supply food chain?

Apart from the illegal aspects of the previous discount arrangements, what may have also triggered the franchise marketing frenzy was the rise of the reseller who is not tied to a supplier. This has been of concern to the incumbent carriers for some time, as these small entities may be the vehicle for future entry of other competitors post the 1997 open marketplace. More competition — lower prices?

The establishment of exclusive franchise resellers is a new development in the cycle of marketing such services. In normal commodity markets it has been observed that sales increase in a linear fashion as the number of general sales outlets increases. But if you make the outlets exclusive sellers of the product or franchisees, then many more sales are achieved for a smaller number of outlets.

It appears that the plan is to establish vertically channels that increase the presence and to offer special prices and packaging that is not available elsewhere. The mobile scene transported to all other services. In some cases this may mean the replacement of the current direct sales arrangement with outside venture reseller organisations that each are majority privately owned franchises with the wholesale carrier having a fifteen percent or so minority stake. The ultimate structural separation without the risk. But can you franchise duopoly rights? Definitely not in the Pay TV world, so why in telecommunications?

Is there any real advantage for the user in the carriers introducing a new raft of resellers between them and the service sources? If the marketplace is fixed then the cost increase is either to the carriers by reducing margins or to the user by increasing costs. Of course, that assumes the direct sales force remains, which is not the case in other industries, but which is unknown in telecommunications. An alternate twist comes when the type of deal is considered that would need to exist between the franchisee and the basic service provider. For the non-dominant Optus the deal could be anything, as the Act allows such freedom, but for Telecom as the dominant provider the rules are a little more stringent, in that the intent was to have all prices available to everyone, with any discount schedules available on the same basis — viz published price lists. There have been some recent minor modifications to accommodate Flexi-Plans, (the discount you have when you're not having one) but the principals have remained.

Now introduce a score of partially owned principal franchises that exclusively resell the service range, with a myriad of sub-agents or spiders underneath. Is Telecom internally transfer pricing or actually reselling against a price list? Can you actually franchise the duopoly? A way round the legislation is to tilt the field!

All franchising appears to do is introduce more aggressive marketing and packaging of the same services, but with competing names that are potentially under less Austel scrutiny than before, and with discounts that are now not open to anyone else. The terms of business would be fascinating with regard to who takes the risk, where the benefit lies, and the credit arrangements. It would be fair to say that without franchising there could be better deals for the users direct, but that is not in the vertical channel game.

Resale of services was one of the early planks into an open market. It was believed then that it would create the necessary innovation and future service providers. The concept looks like it has the legs to introduce new services but if the current moves are successful then resale may be stretched to significantly lower the service prices. I can't remember when I last saw a discount Big Mac.

Tom Amos is a partner with telecommunications consulting engineers Amos Aked Swift.

Kevin Morgan

It's Time to Privatise Telecom

Some months before the last election the then Shadow Minister for Communications, Warwick Smith, argued that the sale of Telecom was not a matter of if, but when. I disagreed, saying that there was no good reason to privatise the national carrier (see 'Privatising For Political Power,' in the October 1992 edition). Now I agree, Telecom must be sold.

I can only say that my earlier views were formed before Austel finally endorsed preselection ballots and before Professor Hilmer launched his pro-competition epistle on the Australian public policy scene. The irrationality of the balloting process and the damage that a national competition policy would do to the publicly-owned carrier have changed my mind; Telecom is no longer safe in the public sector. And if I had any doubts about the need for the national carrier to seek protection in the private sector, where a powerful constituency of shareholders could argue for its interest, they were finally removed by the threatened intrusion of the industry regulator Austel into the proper functioning of an efficient telecommunications market through its so-called Decision Making Framework. If one can disentangle logic from this new Austel approach to pricing behaviour then Telecom has a bleak future indeed.

Austel's new approach to pricing has Hilmer as its underpinnings and it sets the scene for the inevitable structural separation of Telecom so that costly and inefficient infrastructure based competition can be sustained post-1997.

Yet whilst I may owe Warwick Smith an apology for doubting his foresight let me stress that unlike Paul on the road to Damascus my conversion is not due to any new found faith. I maintain that in an ideal world there is no compelling case to sell Telecom. I still believe there is no proof that private ownership results in greater efficiency or heightened responsiveness to the customer. Regrettably though this is not an ideal world. It is one in which pro-competition ideologues set the policy agenda and in which Telecom will ultimately be broken up to sustain competition.

Given the further damage that slavish adherence to pro-competition policies will do to Telecom, my conversion is as expedient as it is pragmatic. Telecom needs to be placed beyond the reach of a shareholder who has little interest in the company's value. The national carrier needs to enjoy a less disadvantaged and less ambiguous relationship with Austel, the regulator which the Government is happy to use as an agent to lessen the company's value. And beyond such expedient considerations lies a longer term benefit from privatisation, access to capital which would allow the marketplace to make judgments about whether Telecom's plans were rational and to the national good.

Also the world has changed. Soon only Telecom and the North Korean telephone company will be left in public ownership! Of the world's top twenty carriers (in which Telecom ranks eighteenth in terms of turnover) the only other publicly-owned carriers are Deutsche Bundespost Telekom and France Telecom. They too will soon be sold.

It could also be argued that in a very real sense Telecom has outgrown public ownership. Building a national network that provides universal coverage and access was best undertaken by the public sector but public ownership it is not essential to maintaining universal service. That could be readily made a licence condition for a privately-owned carrier which would be able to continue cross subsidisation of rural and remote services. Internal cross subsidisation of such services within a privately-owned carrier would be far more efficient and equitable than the web of subsidies, taxes and levies that would be needed to maintain service post-1997 if Telecom were broken up.

And in an age of 'globalisation,' with the universal service mandate discharged, public ownership is not appropriate to the risks implicit in a commercialised national carrier. Indeed, the very nature of Telecom's commercialisation has created irreconcilable tensions. As shareholder, the Government has little wish to see the value of Telecom increase. Indeed its whole policy of competition is predicated upon transferring market share and consequently value to Optus to create so called 'commercially' sustainable competition. With the shareholder behaving in this way one can appreciate the difficulties faced by the Telstra board who are caught in the unusual, if not unique position of answering to shareholders who wish to see the value of the company decline! The directors face an acute dilemma. They see value being stripped from Telecom through such vehicles as subsidised interconnect, and must accept this irrationality because it is the shareholder's wish.

Given the shareholder's attitudes, the board is under no obligation to pressure a management team or employees to maximise the value of Telecom. Consequently they are prepared to tolerate a culture within Telecom which delights in arcane economic debates with the regulator and Optus over interconnect prices. They apparently do not, as BT does, insist on a hard-nosed commercial approach to interconnect pricing. It is this culture within Telecom that has led to a subsequent interconnect price of some 3.5 cents per minute being deemed acceptable despite the fact that this interconnect rate is half that paid in the UK, Japanese and US markets.

Nor does a culture dominated by the shareholder's short term political horizon engender the long term thinking needed for success in the telecommunications industry, where payback periods for investment typically far exceed the three-year term of a Federal Government. And the Labour Party, singed by the failures of entrepreneurial State Banks may understandably be nervous as Telstra's focus shifts from the guaranteed returns of the basic telecommunications network to more speculative investments in overseas markets and broadband cable where the returns are as yet unclear.

Public ownership was absolutely right in the age of basic network extension but for two prime reasons it becomes a distinct liability in the age of competition. First and most significantly, Telecom is the prime victim of the Government's enthusiasm for competition and secondly, in the competitive international marketplace

The Top Twenty Telecommunications Companies

RANK	OPERATOR	COUNTRY	REVENUE TOTAL ($US M)	REVENUE CHANGE 1991-92 (%)	MAIN LINES TOTAL (%)	MAIN LINES CHANGE 1991-92 (%)	EMPLOYEES TOTAL ('000)	EMPLOYEES CHANGE 1991-92 (%)	OWNERSHIP
1	NTT	Japan	51,354	1.7%	57,300	2.7%	242.3	-6.0%	PP
2	AT&T	United States	39,580	2.0%	n.a.	n.a.	77.9	-0.5%	FP
3	DBP Telekom	Germany	34,550	14.3%	35,421	5.6%	231.0	0.9%	SOE
4	BT	United Kingdom	23,379	-0.7%	26,084	1.9%	170.7	-18.9%	FP
5	France Telecom	France	23,164	5.9%	30,100	3.4%	155.3	-0.5%	SOE
6	SIP	Italy	17,492	10.8%	23,709	2.8%	87.5	-0.1%	PP
7	BellSouth	United States	15,202	5.2%	18,677	3.4%	97.1	1.1%	FP
8	Nynex	United States	13,155	-0.6%	15,699	1.9%	70.9	-1.8%	FP
9	GTE	United States	12,644	0.6%	16,819	3.6%	81.1	-4.7%	FP
10	Bell Atlantic	United States	12,093	2.5%	18,181	2.4%	71.4	-7.2%	FP
11	Telefonica	Spain	11,279	14.5	13,792	4.0%	74.4	-1.4%	PP
12	Ameritech	United States	11,153	3.1%	17,001	2.5%	71.3	-3.6%	FP
13	MCI	United States	10,562	11.3%	n.a	n.a.	31.0	11.2%	FP
14	US West	United States	10,281	-2.8%	13,345	3.2%	63.7	-3.2%	FP
15	SW Bell	United States	10,015	7.3%	12,803	3.3%	59.5	-2.8%	FP
16	PacTel	United States	9,935	0.4%	14,551	2.0%	61.3	-1.4%	FP
17	Sprint	United States	9,230	5.1%	n.a.	n.a.	43.4	0.5%	FP
18	Telstra	Australia	8,992	28.3%	8,257	2.6%	70.8	-12.7%	SOE
19	Bell Canada	Canada	6,505	1.7%	9,229	2.3%	52.9	-3.2%	FP
20	Telmex	Mexico	6,185	15.8%	6,546	12.1%	48.9	-1.1%	PP

Note: Table shows major domestic units rather than ultimate owners to reduce double counting
FP = Fully Private PP = Partially Private SOE = State Owned Enterprise N/A = Not Applicable
Source: ITU/BDT Telecommunication Indicator Database

Telstra needs freedoms if it is to emerge as a regional operator. These international ambitions cannot be realised if the Government's enthusiasm for competition goes unchecked and it is questionable whether the current rate of network modernisation and service innovation can be sustained domestically because competition under the Hilmer/Austel DMF model threatens not just the cream for future investment but the very revenues which sustain a modern integrated national telecommunications network.

The Damage Caused by Competition

Competition is one of the twin legs of the OECD/World Bank agenda for microeconomic reform. It's the new orthodoxy and paradoxically, the only escape route from the inevitable damage that more competition will do is the other leg of this reform agenda, privatisation.

In the recent round of debates about a possible Telecom sale the Communications Minister, Michael Lee, dismissed the need to privatise Telecom by claiming in a proud, if somewhat naive boast that Australia had the world's most competitive telecommunications market (whatever that might be) and it was competition rather than privatisation that delivered efficiency gains to consumers.

The OECD and the World Bank of course claim that privatisation and competition are complementary policies. These organisations argue that the correct competitive settings must be achieved in the marketplace before the former public-owned monopoly is unleashed into the private sector. Yet despite such claims about the complementary nature of the two, it is obvious that there are almost irreconcilable tensions between privatisation and competition. Competition diminishes the value of the former monopoly and despite all the rhetoric about getting the competitive mix right before selling, privatising governments ultimately set the regulatory regime to achieve the highest sale price.

Competition has already done large damage to the Australian telecommunications industry and the only bait that might deter the Government from doing more 'competitive' damage is the bag of dollars that would come from selling Telecom as an integrated carrier.

The Phoney Price War

Of course the pro-competition ideologues who think that simplistic neoclassical models of competitive behaviour can be applied to an industry where engineering economics have reached an art form, claim that competition has brought us benefits. They claim that more competition will bring us even greater benefits because in the words of Austel's oft repeated Economics I essay 'competition enhances allocative, dynamic and technical efficiency.' In more readily understood terms it is claimed that competition has led to savings of 20% for consumers on STD calls.

Well, I for one wish to be saved from competition's benefits because as a nation I know we can't afford them. Regrettably, the competitive reality lies far from the world of theory that Austel and like minds comfort themselves with. The reality in the Australian marketplace is no different to that found in the US which former US Department of Justice official Peter Huber has described as 'unnatural' competition. In a 1993 study Huber outlined the reality of competitive behaviour in America. The reality is AT&T sets a price and MCI/Sprint take a price without any real competition between the big three. Huber notes that if real competition were permitted

AT&T would rapidly wipe out the others because the engineering economics of fibre, which scream natural monopoly, have overtaken 'elegant economic theory.'

The same non-competitive umbrella pricing can be seen in Australia. Telecom has set STD prices which are only marginally below those of the pre-competition era. On the published tariff they are only 8% below pre-competitive prices. Optus pitched its tariff over 18 months ago a further 8% beneath the Telecom price and despite the weekend sales and holiday specials the published prices have not changed. Telecom is the price setter and Optus the price taker.

Consequently in a thick route STD market worth some $1.5 billion a year competition offers consumers theoretical savings of up to $240 million a year. The real savings have been far below this and might generously lie around $100 million a year because technology would have driven prices down irrespective of competition, as it did in the early 1980s.

And one should not forget that competition is not without costs — namely the costs of the capital and operating costs Optus incurs to duplicate Telecom's. These costs amount conservatively to $400 million a year! With the costs of competition exceeding the real savings by a factor of four the irrationality of competition is obvious even without delving into the three GSM networks which are rumoured to only have 6,000 customers!

Given these facts Optus has not sought a price war because it knows it would rapidly lead to a Compass Airlines-style exit from the market. Optus neither seeks nor creates competition. But the real harm that competition has done is best revealed by the longer term changes in the Telecom STD tariff. Telecom's STD prices did not change for the four years prior to the introduction of competition despite the fact that in the previous eight years, from 1980 to 1988, STD rates fell 50%.

Quite clearly once the 1988 review signalled the inevitability of network competition, Telecom stopped passing on to consumers the efficiency gains that flow from new technology. Telecom sought to conserve its prices for the reality of competition. The really curious feature of this behaviour is that in the first half on the 1980s prices fell dramatically despite the fact that Telecom's labour force peaked at some 94,000 employees in 1986 and network extension into rural and remote areas was proving to be costly. Yet in the period 1988/1992 when Telecom slashed costs by ditching some 25,000 employees the consumer saw little if any benefit!

Telecom need not have conserved its prices in that earlier period because the competitive reality has proven to be far less threatening to Telecom than feared. The reality is that Optus exists not because of its superior efficiency, better work practices or more advanced technology. Optus exists because Telecom has left it a generous margin in the STD market and price caps now prevent Telecom from closing that margin.

What is truly irrational about this supposedly beneficial and efficiency-inducing competition is that it's focused on the long distance and international segments of the market. The international market is readily attacked because of the price fixing cartel that has operated under international accounting rates. Optus does not need to be efficient because it can merely cherry-pick within that price fixing arrangement. In the STD market, where Telecom has achieved international best practice, Optus merely replicates Telecom's infrastructure despite the fact that Telecom has an oversupply of fibre.

The long distance segment of the market represents only 15% of the assets in the national network and the irrationality of so-called competition is that it encourages uneconomic entry in 15% of the network where prices are above cost and efficiency was already high. Because of the price capping regime the policy precludes competition in the local network where prices are demonstrably 'inefficient' in economists' terms. Any move by Telecom to now introduce more efficient and rational STD pricing is naturally deemed to be anti-competitive and it is this fear that Telecom might be engaged in predatory action that required Austel to grapple with its Decision Making Framework.

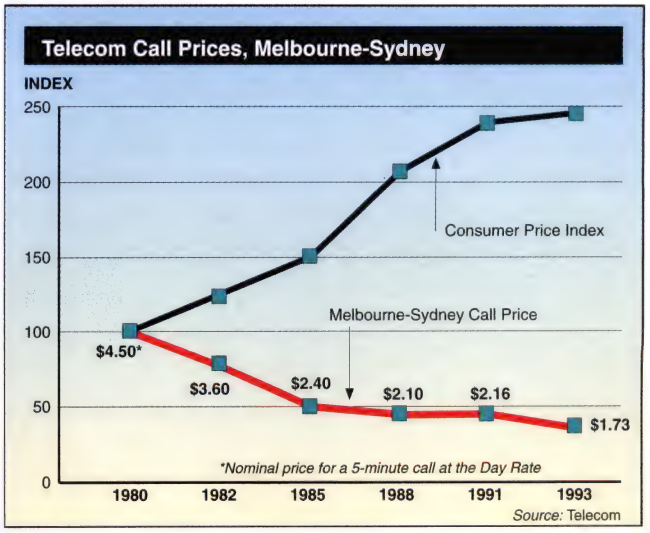

Let's Unblock the Austel Bottleneck

Whilst one can sympathise with Austel for having been asked to come up with the impossible — rational competitive pricing rules in an industry where the underlying economics deny competition — one cannot be tolerant of its initial approach which will turn competition from being a costly irritant into a ruinous disaster. It is this threatened disaster under the DMF with its implicit threat of structural separation that makes it imperative that Telecom be rescued from competition.

In late June, when Austel unveiled the DMF, it argued that predatory pricing could be identified if a price for a service was below the price of the constituent parts of that service. The price of these constituent elements has to be 'flagged.' Under the Austel proposal, which seems to be a curious mix of FCC's mid-1980s access charge model overlaid with the aborted unbundling exercise of Basic Service Elements, the carriers are obliged to offer these constituent parts (so called Basic Carriage Services) to other carriers and service providers at the price they charge themselves.

Without seeking to unravel the underlying logic it would appear that Austel have twigged to the fact that the way to preclude predatory pricing is to effectively bar price competition! This might seem an unduly jaundiced view of Austel's approach, but for the fact that it is underpinned by a curious application of the mid-1980s US concept of 'bottleneck power.' This is a concept which defies any understanding in the real world of delivering telephone service.

Austel reasoned that there is no bottleneck power on the local loop for an originating call because the customer has a choice of carriers and implicitly a choice of transmission medium, but strangely at the terminating end there *is* bottleneck power. One presupposes that the origin of the problematic terminating bottleneck power is the fact that each call is addressed to a unique number, owned by a particular carrier. This fact is of course very useful for users of the telephone network because the call is delivered to the number they have called and not some random number! I am sure users like this engineering reality even if it does induce an unfortunate economic phenomena which might lead to anti-competitive behaviour!

It would seem, given this curious logic, that Austel do not believe that having multiple carriers produces competition. I will agree with them on that. Yet whilst I might welcome their implied conclusion I do not warm to the promise that Austel hold for conduct in the post-1997 market. Under the Austel/Hilmer approach structural separation is inevitable and all incentives to invest in infrastructure will disappear.

Given this threat I wish to be spared the expense of more competition, even if the price is selling my soul and selling Telecom!

Kevin Morgan is an independent consultant based in Melbourne, VIC.

Who is an *Australian Communications* Reader?

- There's a better than even chance that they work for a company with over 500 employees
- There's a 70% chance that they authorise or specify the purchase of data networking products
- It's even money that they run a Novell network or will purchase one in the next 12 months
- There's a 70% chance that they are running Windows or will be in the next 12 months
- And there's an almost 80% chance that they prefer *Australian Communications* for depth of information and value to them in their jobs over any other Australian IT publication

And there are 26,500 of them who read *Australian Communications* every month. Maybe you should be talking to them!

***All statistics are based on the response to the August 1993 *Australian Communications* readership survey conducted by Price Waterhouse, and *Australian Communications'* Audited Paid Circulation.**

For a copy of the complete survey results please contact Craig Burkill or Dawn Switzer on (02) 264 2200.

Price Waterhouse

Dial M For Movie: Video-on-Demand

Video-on-demand promises to deliver Schwarzenegger to your sitting room and Bogart to your bedroom — whenever you so choose. It's a good idea, but does anyone know how to do it?

There's a lot of rubbish being talked about video-on-demand (VOD) — also called 'video dialtone.' The terms clearly mean that individual businesses or households can signal to some central location and a video of their choice will immediately be played down a channel for their exclusive use. It implies a one-to-one relationship between the 'source' and the 'sink.' It's the same as going to the video store and hiring a cassette for your home VCR, but without the journey.

Unfortunately, in the Australian confusion about all this 'superhighway' stuff, the guidebooks on terminology aren't being consulted. Video-on-demand is constantly being confused with:

- *Basic Subscription TV* — a monthly payment made for a good quality electrical signal; for the 'carriage' not the 'content;'
- *Pay (Premium) TV* — a monthly-subscription service for premium programming;
- *Pay-Per-View (PPV)* — which is Pay TV with hourly, rather than monthly charges. You usually pay over the telephone by credit card;
- *'Impulse PPV'* — the newer form of PPV depending on a high level of interactivity over the coaxial cable, often using smart-card authentication; and
- *Near-Video-On-Demand* — which is a Pay TV multi-cast system with time-staggered programming on a group of channels.

These are all related-but-different concepts. None of the above has any one-to-one relationship between source and sink, and all of these techniques present quite different problems to both the carrier and the content provider.

Video-on-demand has recently begun to muddy the waters even more: VOD is conceptually nothing more than database access and information retrieval. The transaction is conducted on a one-to-one basis, computer-to-computer, and the data is delivered as a 90 minute interleaved voice/image 2Mbps digital stream. The problem is in providing the server access for 30,000 users wishing to view 10,000 different movies and 30,000 switched channels to the homes.

Near-video-on-demand (NVOD), refers to the system where home viewers can request and receive access to a small range of popular movies — but on a one-to-many basis — and often with a time delay before delivery. NVOD can be analogue or digital — although digital is obviously better. The server side is easy; the problem is in providing enough analogue channels.

In some ways VOD and NVOD are therefore rival approaches to entertainment delivery — but from another perspective they can be seen as complimentary. NVOD fills a gap between the mass-media requirements of general broadcasting, and one-to-one requirements of VOD viewing.

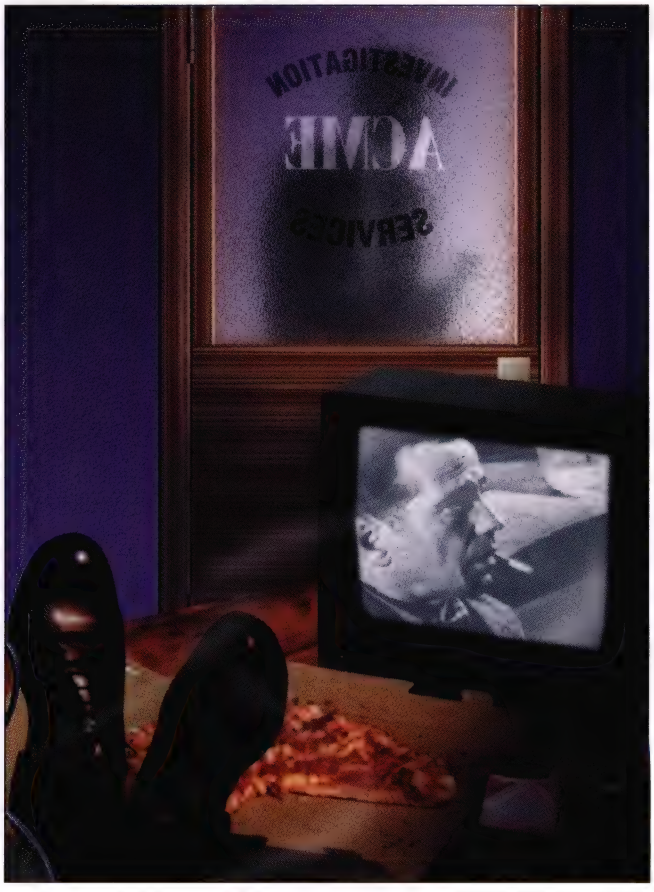

Since, on any one night, thousands of couch-potatoes will be willing to pay for the latest Schwarzenegger release, the program replay devices can be allocated six communications channels, with each replay stream staggered in time by fifteen minutes. So for the one movie they need six 'wholesale' channels (from the national hub to the various head-ends), and each access 'feeder' network

Some Different Ways Video is Demanded

In the US, Cox Cable is testing the Fibre Service Area concept just outside San Diego, California. This uses a star-bus architecture with single active fibres to local serving areas, which then distribute to between 500 and 2,000 homes over coaxial cable. This technique guarantees that there are never more than four active amplifiers in the coaxial line between the head-end and the home.

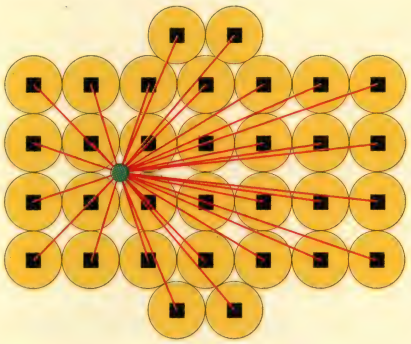

Cox divides the cable spectrum in a reasonably conventional way, providing downstream delivery to the homes between 54MHz and 750MHz, and an upstream return band from 5MHz to 42MHz. But it then subdivides these bands to provide digital telephone (including PCS wireless) and multimedia/on-demand services. Currently all video delivery is analogue, but this will be converted to digital soon.

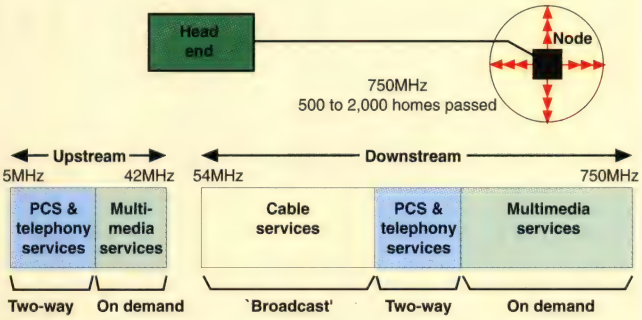

A conventional coaxial cable network has a tree-branch architecture consisting of head-end amplifiers with multiple trunks running down the major streets using 0.75-inch coax. Distribution splits using 0.50-inch coax carry the signal into side streets, and passive splitters then send the signal to individual homes. Bandwidth is divided between upstream and downstream functions. Typically, from 5MHz to 30MHz (recently 42MHz) is available for the upstream return channel, and 54MHz to 550MHz (or more recently 750MHz) for video. The upstream channel has been widely used in the UK for voice telephony. These systems may have 30-40 amplifiers between the head-end and the home, and so there are constant problems of reliability.

Time-Warner has a hybrid fibre-coaxial service already operating commercially in the Queens district in New Jersey. This uses a conventional cable topology, but the main coaxial trunks have been replaced by a bundle of fibres to reduce the reliance on trunk amplifiers, improve reliability and picture quality, and provide a greater number of channels. This is an all-analogue system providing 150 NTSC channels between 54MHz and 1GHz, and messaging for Impulse Pay-Per-View over the 5MHz to 35MHz upstream band. The upstream band and a sub-section of the downstream bandwidth will be converted to digital in the near future.

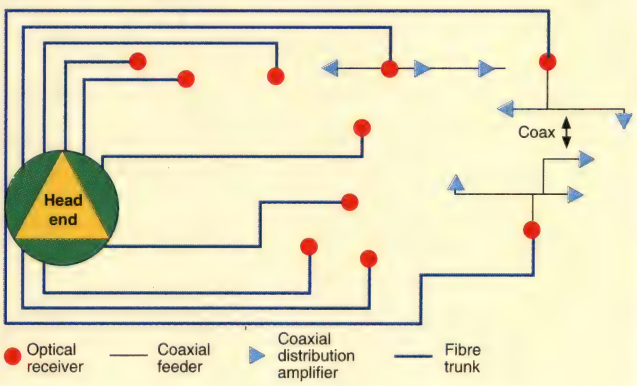

Time-Warner's Orlando (Florida) trial is for a Full Service Network which includes video, telephony and wireless links. It is a hybrid analogue-digital system using MPEG-2 compression and ATM for on-demand services. The conventional analogue channels are merged with the digital video-on-demand streams at Digital Node Controllers (DNCs). The signals are then modulated into 60 DS-3 64-level QAM (12MHz wide) channels and sent over fibre to local nodes which service about 400 homes over coaxial cable. PCS (CDMA) cellular base stations exist at these node sites for low bandwidth telephony and data services, and separate umbrella cellular base stations provide wide area coverage.

The 1GHz of available bandwidth is divided to provide video-on-demand, electronic shopping, navigation systems, interactive games, and telecommunications. The important telecommunications component includes switched local services (entirely within the cable network), alternative access (to conventional telephone services) and mobile communications services (wireless used for PCN, cellular mobile, telepoint and other cordless applications).

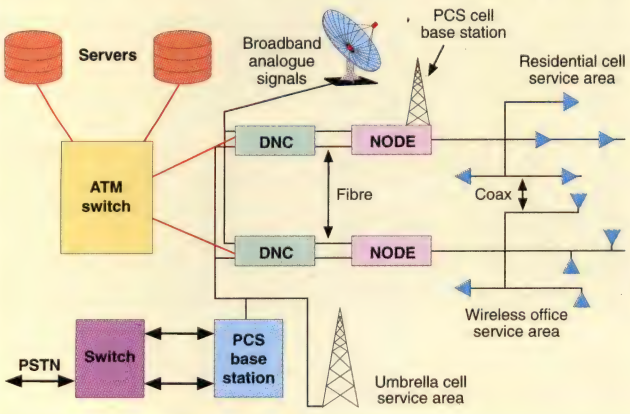

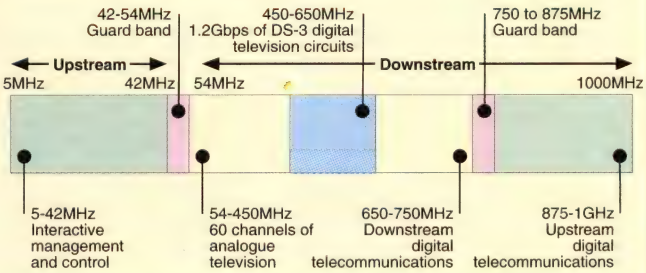

also needs to have six channels available. They can source these time-staggered images from a half-dozen replay machines (VCRs, video disks, hard-disk or RAM-disks) with the same movie, or by swinging a number of heads over the surface of a single disk — each reading a different part of the disk.

However, the obvious technique to use with modern high-speed Pentium and Power-PCs equipped with PCI (Peripheral Component Interconnect) local buses (which can

handle 6 streams of 2Mbps without blinking), is to source all program streams from the one disk, relying only on fast-disk access and buffering for the time-staggering. In the future, the most popular movies will probably be held in large RAM-disks simply to save on mechanical wear and tear.

There's no complex video server operation needed here: the problem in providing such a multi-cast service lies in the economics of being able to allocate the six channels to one movie. And there is no necessary relationship between time-staggering and the payment mechanism.

NVOD movies are usually charged on a pay-per-view basis using telephone-request or smartcard billing systems. But this isn't the only way time-staggered channels can be used. A cable company could well offer a few dozen monthly subscription channels with staggered blockbuster programming, and it is likely that free-to-air (advertising supported) television broadcasting will use NVOD time-shifts when there were plenty of digital channels available just to make viewing more convenient.

Obviously it is easier to provide 50 NVOD channels for shared viewing than 5,000 one-to-one channels for true VOD. So the first requirement for is to have enough broadband channels (7MHz for analogue, and 2-4Mbps for digital) available through the main 'wholesale' network (to head-ends, exchanges, or transmitters), and then out through the access network, to the viewers.

Cable and satellite expert, Robert Cooper, dimensions the two different problems in a recent *Coop's Technical Bulletin*. He points out that:

- More than 35% of video store income in any month comes typically from the rental of fewer than ten films; and
- Half of all rental-store income comes from the top 21 films.

So these 21 films are prime candidates for NVOD — while true VOD will be needed to look after thousands of random movie titles chosen by the other half of the viewing audience. According to *Fortune* magazine, the largest US video rental stores carry 10,000 titles which, on disk with MPEG-2 compression, need the equivalent of 20,000 gigabytes of storage. That's 4,000 five gigabyte hard-disk drives, folks! That's a large warehouse at least, along with dozens of maintenance technicians.

Cooper's statistics also put a perspective on the relative complexities of the two problems. If they have 60 channels churning out the top ten films staggered at fifteen minute intervals, NVOD should satisfy about one third of the potential audience. With 99 channels and half-hour staggering of the 'next 11' they can satisfy half in total — assuming that the popularity of electronic delivery systems roughly parallels that of store rentals (which, of course it won't).

Who Carries What?

In the US, telephone and cable TV operators have long been prohibited from entering each other's territory, which is why the use of shared-cable for telephony has never developed. This is now changing slowly, state by state.

By 'integrating' the 'broadband services' over radiating optical fibre (Broadband ISDN) the world's telcos thought they could lock up the provision of information, and squeeze out the hated cable companies. But in the UK and Europe, a new open environment allowed cable companies to offer local telephony links, and they then began to provide long-distance services through deals with BT's rival Mercury.

The cable companies had discovered that telephone links could be carried back to the 'head-end' (the primary amplifier racks and cable trunk distribution point) using the 5MHz to 30MHz part of the cable spectrum. So the cable companies 'integrated' back!

An average cable system will reach (pass) between 10,000 and 20,000 homes from a single head-end (the main feed amplifiers), but each cable-trunk may only feed between 1,000 and 2,000 homes at the most. So even if a standard frequency-division system were to be used (which no one does) it would still be possible to get 1000 phone circuits (at 4kHz each) into 4MHz of bandwidth on each trunk — and the upstream link has 54MHz available.

The telephone companies have, for years, believed that they held the upper hand by being able to introduce broadband fibre to the home in a star-network, and thereby control multimedia delivery via their telephone exchange 'hubs.' ADSL over twisted-pair was just another attempt at locking down these services, but without the expense of opto-electronics. But ADSL is now so expensive that it increasingly looks like a niche-market technology.

Cable companies in the United Kingdom (some owned by US telcos!) now offer discounts of 50 and 60% for telephone connections, and have taken more than half the local 'dialtone' market away from British Telecom in many areas. BT sees this as a serious long-term crisis, not just a temporary anomaly.

Stewart Fist

I don't accept the 10,000 hard-disk concept, because we are many years from being able to devote storage space on hard-disk replay units to obscure Laurel & Hardy movies which may be only requested once a month. But there is an alternative approach for the less popular movies: we can use a silo system where a machine goes away and selects the cassette or optical disk you want and loads it into a VCR/CD-I player. It is still video-on-demand — just not instantaneous.

In both forms, the channel requirements are the same (one per using household), but a silo system servicing 5,000 households would probably not need much more processing power than a 486 PC since it processes the preliminary request/selection only. By contrast, the instantaneous approach calls for supercomputer-power at least, since the control system must be continuously accessing 5,000 bit-streams emanating from hundreds of hard/optical-disk units.

Instantaneous video-server operations require a new order of computer complexity. To take an example, American Airlines' Sabre reservation system, used by airlines around the world, is one of the world's largest transaction databases — yet it needs only 100 gigabytes of disk storage. The point is that it has enough VOD capacity for only 50 movies. So to match the local video store in variety, each major city 'domain' will need the equivalent of at least a Sabre system in transaction processing capability, but with access to about 100 to 200 times the storage capacity. This is a tall order.

No matter how powerful the modern microprocessors, or how massively parallel they make the central processing systems — or how densely they cram bits onto magnetic or optical disks — downsizing 200 Sabre systems to a convenient, cheap, commercial video-server unit is still not a trivial matter. I would suggest that for the rest of this century, local instant-access storage will only be possible for, say, the thousand most popular films. Silo systems will be necessary for the other 9,000 titles and we'll get very used to messages like 'Please wait while I load a cassette.'

Coax Crisis?

Modern analogue coaxial cable networks carry 96 channels of PAL (at 7MHz spacings), and about 110 channels of NTSC (at 6MHz). So cable companies can provide a basic NVOD service using, say, 60 analogue channels, and still have 30 or 40 left for conventional basic or Pay TV. The bandwidth of these cable systems will increase even more when fibre begins to replace coaxial trunk feeds in about two years' time.

Time-Warner's trial hybrid network in the Queens borough of New Jersey is a good example of these new cable networks. It currently provides 150 analogue channels by reducing the number of trunk amplifiers to a maximum of four (see figure) between source and sink. Conventional full-coaxial

The Super High Frequency Future

Bandwidth has been a key concern of the analogue era, so when optical fibre to the home held out promise of infinite channels, terrestrial radio systems and coaxial were discounted. But that's now changed. There is plenty of room for a few thousand analogue channels (3,857 to be exact) up in the Super High Frequency region (3-GHz to 30GHz) if we wanted to use the whole spectrum band in this way.

SHF is the next band above UHF, and it extends across a massive bandwidth of 27 GigaHertz — which, you will notice, is nine times greater than the combined total of all the radio frequencies we have used to date. And since bandwidth is related to information-carrying capacity, SHF promises video channel space to burn.

As with all historical efforts at extending the useful spectrum range, until recently they've had problems building commercial receivers able to operate at the new band of higher frequencies. There were also doubts about SHF's ability to penetrate more than a few kilometres of atmosphere, especially in rain or fog. Five to ten kilometres seems to be about the practical limit, with one or two kilometres being much better.

CellularVision is now trialling such a system in Brighton Beach near Brooklyn, New Jersey — and these trials demonstrate how fast radio technology is now advancing. In CellularVision's SHF trials, 2GHz at the *top end* of the band (in between 26.5GHz and 28.5GHz) is being used to provide 50 channels of NTSC 'Basic Cable' and some 'Premium' channels. They 'spread' each channel across 20MHz of spectrum, which wastes analogue space, but not digital.

Near-video-on-demand channels will be added shortly, along with a radio-linked interactive component probably carrying voice-telephony as well as messages and billing/selection information. Digital transmission is still a couple of years away for all terrestrial radio systems.

So at present, CellularVision is little more than MDS (Microwave/Multipoint Distribution System) translated from 1.6-GHz (in the UHF band) to 26.5GHz — but it has proved the concept that SHF works, and works well. In doing so, it has re-established terrestrial broadcasting as a superhighway technology, and put it on the same footing as satellite and cable in terms of large channel numbers — in the 150 (analogue) to 1,000 (digital) region.

CellularVision has also proved its financial and technical viability to the point where Bell Atlantic and Philips recently bought half the company. And in the hurricane which hit New York last year, SHF continued to operate without a break while most of the surrounding cable systems were demolished — mainly because US coaxial cables are strung on poles.

The SHF band has similar restrictions to MDS in needing line-of-sight to the mast-head, but MDS only provides us with 19 analogue channels while SHF is open-ended. But, as with cellular telephones, these limited-range restrictions on the signal are turned to advantages as soon as you abandon the idea of one-way broadcasting to the masses. The smaller the cells, the more the diversity.

SHF's line-of-sight transmissions require the use of small cell structures — usually a transmission distance of two kilometres maximum. But the trade-off is that the system operates with low-power, provides hundreds of channels, and can be received by very small dish (or phased-array) antennas.

The benefits of the short transmission range are substantial; there's little in the way of multipath (ghosting) problems, and the received-power of all channels is matched since they all emanate from the same source/head-end. This is a very important point that is often not very well understood.

The real bonus comes from the fact that a home with a transmitting set-top box can easily reach back to the head-end with messaging, billing information and voice telephony — it can double as 'radio in the local-loop.' So SHF is likely to challenge both cable and satellite networks for VOD, and since it can carry voice telephony, it may even affect the local-loop monopoly of the telephone carriers.

Stewart Fist

systems may have 40 trunk amplifiers along the line (from the head-end) before the signals reach a TV set.

Hughes DirecTv satellite system (which is digital) claims to provide about the same number of channels as the Queens trial from two satellites. Satellites have a greater need for digital compression than cable systems at the present time, and it is not likely to change in the future. Generally there's about a 4-to-1 ratio between the number coaxial can carry and what satellites can provide for DBS, and the ratio is shifting more in favour of cable every day.

With a Queens-type hybrid network the 21 most popular films can be provided as NVOD easily, even though they are still using analogue transmission. On average, the Queens viewer waits less than ten minutes to see his/her selection, with a half-hour worst-case. That's still quicker than going to the local rental store.

Currently about 8 million North American homes and 25 million European homes watch PPV NVOD services via satellite or cable (with an average 15 minute staggering), and experience shows that this doubles the service provider's income over Subscription Pay services. In New Zealand, the average home orders between 5 and 6 movies a month, which roughly matches US and European statistics.

Until digital compression and modulation comes into the picture, this is the state of the art for on-demand systems. NVOD is currently feasible with analogue TV, while true one-to-one VOD is not, because of the limited number of channels available.

True VOD over analogue access networks is only possible with a switched network carrying only one signal in the local loop. Unfortunately, our current access and inter-exchange network are both incapable of handling analogue video bandwidths.

Shared or Switched Media?

The real driving force behind the 'information superhighway' concept is entertainment not information, so let's stop pretending that we have more worthy motives. I am heartily sick of attending Pay TV conferences where potential programmers preach pious pap about channels filled with opera, ballet, museum/art gallery tours, and the like. Sure, these will come, but only once the sport, movies, cartoons and bikini mud-wrestling appetites of the audience are satisfied. The Barbarians are at the gate.

This is where the pop-culture concept of '500 channels' comes in. I thought you'd like to know that the 500 figure was conjured out of the air during a TV interview by Bell Atlantic's marketing manager — and he was talking about ADSL (Asymmetrical Digital Subscriber Line, which can provide video over twisted pair wiring; see 'ADSL — Bridging the Superhighway Gap?' in the May edition). But the figure makes no sense. If video is truly available on demand over a 'switched' technology, then you can't put *any* figure on the number of channels available. He could equally have said 500 million.

Amazingly, the futuristic fibre-to-the-home superhighway now appears to be even further over the horizon than a year ago. Some people (including me) are beginning to question whether fibre is necessary at all. If you can access 500 million channels over a phone line, why do you need glass fibre through the streets?

But the real action is in coaxial. The world has discovered that shared coaxial media (as with LANs and cable TV net-

CUSTOMER SUPPORT

EVERY PIECE OF OUR EQUIPMENT COMES WITH OVER 250 MOVING PARTS.

At Nortel, we know the importance of customer support.
That's why we have over 250 fully trained staff providing a range of marketing,
logistics and technical support services to our customers.
Each one has access to worldwide databases and software support centres - 24 hours a day, 365 days a year.
Ensuring that the service our customers receive is of the same standard as the equipment itself. Worldclass.
After all, they wouldn't expect anything less.
Nortel. Discovering and delivering the world's best solutions in telecommunications.

nt nortel

works) can comfortably carry one-to-one and one-to-few services, as well as one-to-many, in truly amazing quantities.

Also peeping out from a cloud of secrecy is a new radio technology called 'CellularVision,' which uses the Super High Frequency (SHF) part of the spectrum. This could also emerging as a fibre-killer; it has all the attributes of coaxial without the cable cost.

Cable and CellularVision subscribers can be addressed individually, or as a group (as with a LAN); they can cheerfully share space between analogue and digital transmission channels; and they can support circuit-, packet-, and/or cell-switching. SHF is still a year or two away, but one of the attractions of rediscovering coaxial networks is that no revolution is needed, just some evolutionary migration of a technology that has been around since 1948.

Interactivity does, of course, make a difference; and it is interactivity on coaxial and SHF that will probably prove to be the satellite-killer. The first interactive coaxial cable networks just used the low-frequencies (below Channel 0 at 40MHz) to carry Pay-per-view request and billing messages back to the head-end. But as digital compression comes into cable, these interactive systems are likely to become much more sophisticated: more LAN/WAN-like, and more telephone-like.

With packet-switching protocols, shared coaxial cable is a much more flexible system than twisted-pair switched-star telephone networks. Coax can provide 'connectionless' messaging, and it can handle bursts of data at rates that ADSL can only dream about.

My bet is that the first direct experience that public will ever have with ATM (Asynchronous Transfer Mode) will be in receiving voice, video and data delivery over their coaxial-copper cable TV connection. The telephone companies may have invented ATM for their futuristic 620Mbps optical fibre B-ISDN networks, but the protocol now serves other masters. ATM works perfectly well at low rates down around 64-Kbps also, and it is genetically endowed for shared-media switching.

So VOD now depends on the combination of MPEG-2 for compression and ATM for switching/selection; with modulation (probably 64-QAM) and transport (multiplexing) techniques still to be standardised for cable and SHF. Theoretically, it should be possible to transmit digital information at about 8 bits for every Hertz of bandwidth available, and the Queens hybrid network has a full 1GHz on offer.

MPEG and ATM may also be used with ADSL over twisted-pair if the carrier is a telephone company — although ADSL doesn't look particularly healthy at the present time. Between 2Mbps and 6Mbps are being claimed for the DMT (Discrete MultiTone) version over distances of 2km from the telephone exchange, and this is enough for a switched individual feed, but the images seem to have a lot of odd artefacts, and MPEG-2 motion compensation does not work too well at lower rates.

By contrast, a Queens-type hybrid network should reach 8Gbps in full flight and that's (theoretically) the equivalent of 4,000 digital movie channels (at 2Mbps) or 2,000 better-quality movie channels (at 4Mbps) or 1,000 sports channels (at 8Mbps), or perhaps 400 High Definition TV channels (at 20Mbps). Whatever bandwidth is required can be delivered.

For VOD over cable, you've got to remember that most of those 2-4,000 movie channels can be made individually addressable down *each* of the access-trunks, provided the 'wholesale' network has the capacity to transport all of these separate signals from the video-server source. ATM will provide the switching to the individual home.

ADSL connections will also need to be switched from a large capacity wholesale delivery system unless local servers are to be provided. We assumed until recently that ADSL was likely to be hobbled by the need for expensive 2-4Mbps circuit-switches on each line at the local telephone exchange level — and this would be a multi-billion dollar upgrade nationwide. But there's no reason why ATM-switches can't be used with ADSL in a data-over-voice way — where frequencies below 4kHz are reserved for POTS, and those from 16kHz to 1MHz are for ADSL/DMT controls and video.

So, the way it looks at present, is that a wholesale network servicing a city can equally link to either ADSL or coaxial access systems. In both cases the normal POTS phones would continue to work with a stock-standard POTS voice switch.

The real question about ADSL is whether the required bandwidth can be delivered over the local loop at a comparable quality to the alternate (coaxial and SHF) approaches — even allowing for the fact that the copper twisted-pair is already in place. Most of the vendor companies only saw ADSL as an interim measure anyway. My guess is that it will find its place serving niche markets in education, information, videoconferencing, etc. — wherever requirements can be satisfied by medium-quality, or reasonably-static video images.

Video Domains

On any shared-media system like MDS/SHF, satellite DBS, or coaxial cable, we'll need very high bandwidths to achieve true VOD — if it catches on. All of the above has assumed a few thousand movies on offer. But add to this educational and training films, and you double the size of the problem instantly. Add to this documentaries and the like, and you double it again.

Then, what about those who want to revel in nostalgia? Those who want old movies, well-loved sit-coms and classic TV dramas? Each week, prime-time TV producers probably generate 200 hours of material that someone, somewhere, wants to see again — and that's 180 Gigabytes at the lowest 2Mbps MPEG-2 rate . . . per week! So where do we stop?

And where are these servers to be located. Do we just have one for Sydney, or will we need one in each suburb? This is the problem of server domains. The parallel is with cellular phone systems in the mobile telephony environment: the limitation on one-to-one contact depends on the number of channels which are available in each 'domain' or cell area. If you want more useful individually addressable channels in the same bandwidth, you need to reduce the domain (cell) size.

Satellite DBS will probably have a single 'broadcast' domain covering most of the populated areas of one State, MDS may have one domain across most of a city or at least covering a number of large suburbs, while SHF, being a small-cell MDS system, has small domains.

Cable systems are able to create variable domains by grouping or separating the main trunk feeds in their head-ends. They can use the whole head-end as a single domain (with, say 10,000 homes) or break it down to individual trunk domains (say 1,000 homes) — either way, the full cable spectrum is available in each.

When hybrid optical/coax cable systems become common, domains will be created as small as a fibre-serving-area — say, 500 houses — which is another reason why fibre will gradually come into these hybrid systems. It is not just to increase overall capacity.

However, while having smaller domains solves the problems of VOD on the access network, it doesn't solve that of the inter-head-end links (the wholesale network) which still needs a separate channel for each VOD user. However the Japanese make 400 fibre inter-exchange cable now, and with each cable shifting data at 2.4Gbps rates, this gives us nearly half a million 2Mbps video channels. So it seems as if the wholesale network shouldn't be a problem.

Video Servers

Even with an infinity of channels, the question remains about video servers. In the current NVOD services, some human slave runs around between a dozen or so VCRs plugging in cassettes, but that's not likely to be feasible with true VOD if a million or so homes are making choices between 10,000 different movies!

Silo systems will obviously handle the more obscure requests, but large amounts of the more popular stuff will need to be compressed onto hard disk or optical disk units, mounted ready for instant replay. The criterion here is going to be the average

"We can sell the Tennis Centre (16,000 capacity) in about 20 minutes."

Manager Information Systems
Bass Victoria Ticketing

Making the business of the box office run smoothly across the State of Victoria is what JNA is achieving for Bass Ticketing.

JNA Network Services gets to know your business first, and your business goals, before extending the network. The technical expertise follows. And that is considerable.

For over 30 years we have implemented complex inter-networking solutions for banks, retail chains and government, city-to-city, building-to-building, floor-to-floor. And we get it right.

So see us first. It is quite likely you need more than a network integrator; you need a business partner. Call us now on: Sydney (02) 417 6177; Canberra (06) 273 5044; Melbourne (03)532 8555; Brisbane (07) 870 3434.

We're building partnerships as well as networks.

JNA Network Services

Innovators in Business Connectivity

capital cost of providing an individual bit stream through a server.

As you'd expect, all of the major computer players are involved in video server design or in developing control software that they hope will become the de-facto standard. There's no obvious need here for international standardisation, so everyone has a chance to win a slice of this potentially lucrative market.

IBM and DEC have taken the approach of trying to repackage their old big-iron, using multiple disk-packs to store the information and depending on their high-powered CPUs to feed it out again as ATM cells. It is generally agreed now that mainframe or mini-based architectures are not the cheapest way to go.

However IBM offers something new in its 'multilevel' optical disc drive, which is similar to recording information on half-a-dozen transparent CD-ROMS, then stacking them together to create a six-layer data sandwich. They then read the pit information by changing the focus point of the read-head, shifting up and down through the layers, at will. IBM says a 10-layer disk can store 6.5 gigabytes (you need about two gigabytes for an MPEG-2 movie). There's a lot of interest here, since these disks can be cheaply manufactured in quantity.

Hewlett-Packard (HP), through its Video Communications Division, is also investing heavily in this race and is developing its video transfer engine architecture. HP is also following the ATM route, but is designing its system from the ground up rather than cobbling together any old hardware that happens to be around.

Alcatel, Nortel and just about everyone else who owns a soldering iron is in the server battle too, most of them with rehashed computer gear which has been souped up to process MPEG-2 video, and video files are being crammed on to the biggest magnetic disks available (usually 5GB). VOD trials are taking place all around the world at present, but most of the systems are servicing only a few hundred to a thousand homes, so these are *feasibility*, rather than commercial, trials.

Apple is using its Power Mac with a PCI local-bus as a server, and it has a large trial in America with full-blown Macs disguised as set-top decoder boxes. Scientific Atlanta and General Instruments are battling it out for other markets; and all the games manufacturers (but particularly 3DO) have partnerships and trials under way. Australia's own AWA has reprogrammed the ABC-invented D-cart sound-editing/replay system, and has produced a system known as V-cart.

Silicon Graphics has also entered this field in a big way (they've got strategic partnerships with just about everyone). Currently they are involved in a number of trials, including the Time-Warner Full-Service Network (FSN) trial with 3,000 homes in Orlando, Florida (see the figure on page 66). FSN is being promoted as 'Cable's answer to the telco's B-ISDN,' but this characterisation refers to service provision (I hope), not cost.

The latest reports say that the Time-Warner consortia is spending $108,000 on average for each home connected — and I'm not sure they are counting the server costs. Silicon Graphics provides a supercomputer as the video server, and currently it can deliver custom-selected movies simultaneously to perhaps 1,000 homes, which is certainly enough for this trial.

However, the two video server projects which stand out at present as having real potential are that of Microsoft/Intel, and the earlier one from Larry Ellison's Oracle. There's been a lot of hype about the Oracle server, because CEO Ellison decided that this was his chance to do a Bill Gates and become famous as the garrulous guru of the golden gateway.

What Ellison figured out was that if a video server is to be capable of handling 1,000 MPEG-compressed hard-disk drive movies and feeding individual data streams to, say, 30,000 homes, the bottleneck was going to be in the video disc controller.

In practical terms, this means that the storage devices must be capable of delivering their data at many times the rate needed for MPEG-2 (let's say 4Mbps for movies). Up to 100 homes may be accessing the same disk (at various stages through a movie) so the bit-stream from each needs to be perhaps 50 to 100 times faster than the individual stream requirement.

For 30,000 homes, the central processor would need to handle hundreds of disks, each sourcing many buffered bit-streams. So the CPU output would need rates of 30,000 x 4Mbps = 120Gbps, which requires a fairly hefty number of MIPS. When Ellison first started to look, this was beyond the capabilities of virtually any machine around: these are modern supercomputer rates.

To cut a long story short, eventually he found the company nCube making a massively-parallel super-processing box with 1,000 microprocessors. The top nCube box costs $30 million, but it runs at 123,000 MIPS — about 200 times faster than the average mainframe. He liked the products so much that he bought the company.

So the Oracle server comprises stacks of standard off-the-shelf 5GB hard-disks (but used in RAID-5 fashion), with the nCube box acting as controller and network manager to pump out the data streams in ATM cell format. Obviously a lot of clever software needed to be written to handle this number of data streams, and they used the Oracle database manager as the foundation. This appears to be the first practical video server solution to escape from the labs.

Oracle hasn't yet reached the figure of 30,000 individual data streams, but the new nCube2 computer with 8,192 microprocessors should do the trick sometime later this year, and next year, the nCube 3 with 65,000 microprocessors should be capable of handling up to 150,000 concurrent users. That's the Oracle line, anyway.

The only thing that worries me about all this is the base-level box costing 'only' $30 million. What is it going to cost me to watch a movie, if I'm one of the few thousand being serviced by this machine?

Luisa Conte of Telecom Research Labs estimates that the capital cost per bitstream figure for nCube will be in the order of $600 (for processing only, not the drives) which is probably tolerable — although it will make a dent in the profits for the first few years if we only watch 4 or 5 VOD movies a month. Don't forget also, that Telecom will charge an arm and a leg for exclusive 90 minute use of a 2Mbps data stream through their fibre-interexchange and co-axial-access network to your home.

Conte points out that the comparable figure for the Microsoft/Intel proposal is only $50 per bitstream. "The difference is not in the processing power — you don't need big processing power," she says. "What you need is the possibility of connecting multiple users to the same movie. Massively parallel systems allow you to spread a movie across different disks in RAID format. If these disks are managed by low-power independent processors in a parallel way, they can make a system which satisfies more users at the same time."

This is one solution, and nCube/Oracle says this is the right way. But Microsoft and Intel have joined forces in developing 'Tiger' which is a hierarchy of PCs running an enhanced Windows NT operating system. Tiger uses the new PCI local bus on special video cards to shift the data to multiple ports. PCI is Intel's bus architecture which is now an open, non-proprietary standard, supported by 160 companies. PCI is processor-independent (Apple uses it on the Power Mac) with a 32-bit architecture, and it boasts a peak bandwidth of 123Mbps. By reducing the bus length, they can step it up to well over 150Mbps — which Microsoft claims is 100 video streams. Personally, using the 4Mbps MPEG-2 measure, I make this 40 independent channels.

The key 'Tiger' machine sits in a supervisory capacity over the top, and a large number of (almost identical) 'Cubs' hang off it and perform the bit-stream processing and disk accesses. According to Microsoft, there is no limit to the Tiger-Cub ratio, but the most they've handled at present is 16 Cubs.

This division of responsibility means that off-the-shelf hardware can be used to challenge the nCube, IBM, DEC and all the other hardware-oriented players.

Stewart Fist is a freelance journalist based in Lindfield, NSW.

Thanks to IDMS, your users never even knew it.

The IDMS monitoring system from Wandel & Goltermann is far and away the best solution for proactive management of your mission-critical LAN segments. That's because IDMS gives you a clear overview of LAN segments, nearby or remote, not a lot of MIB jibberish. The powerful easy-to-use IDMS Manager runs on UNIX management workstations, and with IDMS Remote Monitors is the tool you need to detect and diagnose problems across an entire internetwork – before your users find out the hard way. The IDMS system provides for complete implementation of the RMON MIB standard, and features the power to expand as your monitoring needs change or grow.

But then, what else would you expect from W&G, developers of the future-proof DA-30 protocol analyzer? For the best in worldwide internetwork monitoring and testing, call W&G today.

Wandel & Goltermann
Electronic Measurement Technology

For more information contact Wandel & Goltermann Pty Ltd, 42 Clarendon Street, South Melbourne, Victoria, 3205, Tel. (03) 690 6700, FAX (03) 690 6750

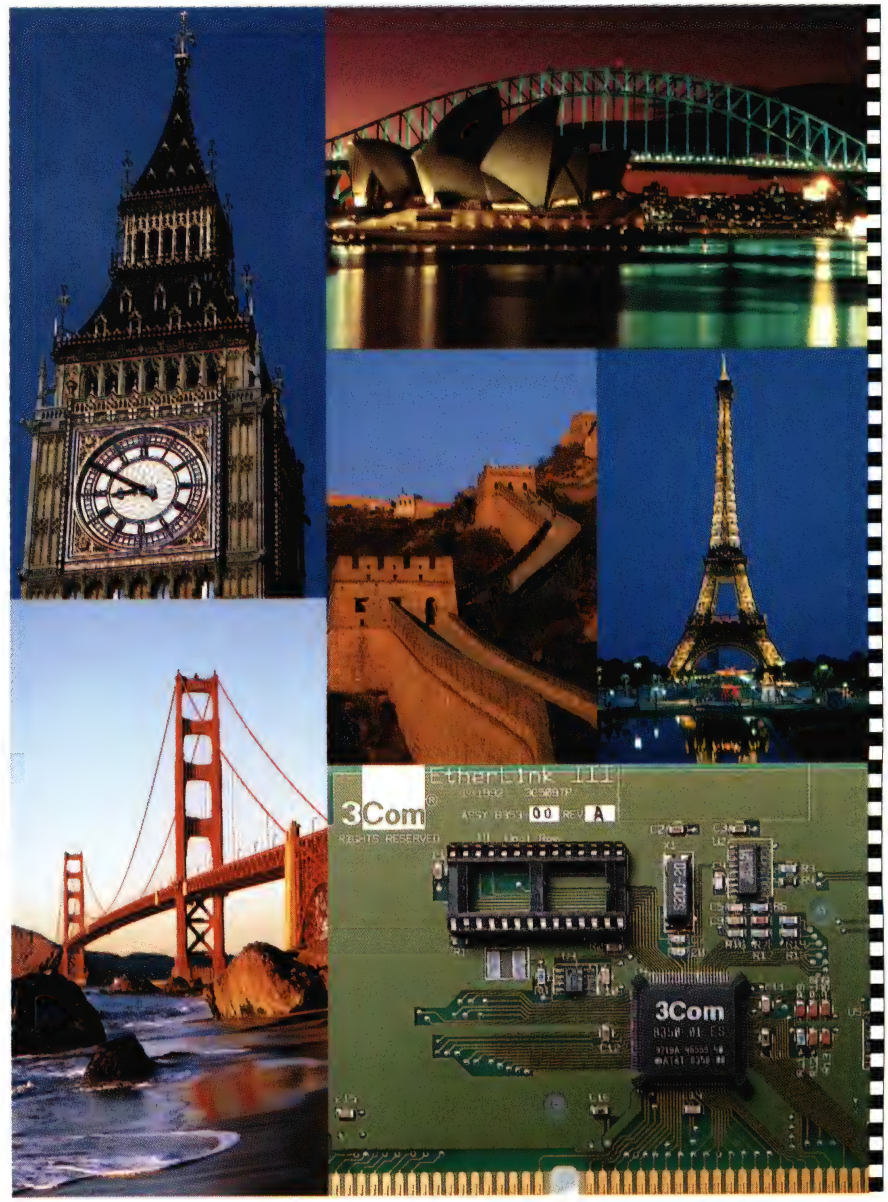

There's only one original Parallel Tasking.™

■ What would you call a technology that delivers the fastest adapters in the world? That has won immediate acceptance from network managers around the world?

Millions know it by the name Parallel Tasking. And you'll find it in only one place: 3Com.®

With 3Com's Parallel Tasking adapters there is simply no faster way to run your network. And there is no other adapter line that offers more convenience and more performance at a lower price.

An EtherLink® III Parallel Tasking adapter is SNMP-manageable. It is auto-configuring. And every 3Com EtherLink III adapter is backed by a lifetime warranty.

It's no surprise, then, that more than two million have been sold.

But why not judge for yourself? Call **(02) 959-3020, ext. 20** for your free *Guide to Benchmarking Network Adapters* from LANQuest Labs, and complete information on 3Com's EtherLink III family of adapters.

It's never been easier to put unparalleled performance on your network.

3Com®

Networks That Go the Distance™

Building Australia's R&D Network

With efficient data networking now recognised as a key element of the information economy, super high-speed research networks are being built in several countries. What should Australia do?

In October last year the Federal Government commissioned the Australian Science and Technology Council (ASTEC) to report on Australia's future requirements for Research Data Networks (RDNs). After issuing some preliminary findings in a draft report released in April, ASTEC is scheduled to shortly table in Federal Parliament the final results of its study, called 'The Global Connection: Future Needs for Research Data Networks in Australia.' During the preparation of its report, ASTEC commissioned Melbourne-based consultancy Cutler & Company to investigate key RDN issues. The article which follows has been adapted from Cutler & Company's findings.

In March this year there were an estimated 2.2 million computers connected to the Internet via 27,000 networks in more than 70 countries. This enormous global net was then being accessed by at least 7 million users, rising to 20 million if those that use the system only for e-mail are included. AARNet (the Australian Academic and Research Network), which is Australia's link to the Internet, currently links around 110,000 computers, each — judging by overseas experience — with four or five users. While it's very hard to judge accurately, traffic entering AARNet (from domestic and overseas users) is averaging 8 megabytes per week per host, or more than 800 gigabytes per week. This is saturating the present 2Mbps inter-city AARNet backbone, and is driving AARNet's present plans to upgrade all major interstate links to Telecom's Fastpac 10Mbps service during the second half of 1994.

But will these upgrade plans be adequate to satisfy demand? AARNet traffic volumes grew by 130% during 1993 and the rate is increasing. AARNet could see 450,000 visible hosts by 1996, and users on those hosts will most likely be using newly-developed user interfaces like Mosaic, which speed and simplify access but consume 30-50K of data per page versus the 1-2K for an ASCII e-mail page. What factors are driving Australia's RDN requirements and what should be done to meet projected short and long term demand?

Applications and Growth

Figure 1 on page 76 sets out a series of RDN applications that require increasing end-to-end bandwidth for acceptable performance, as perceived by the user. The first four groups, requiring RDN performance for the user's payload at speeds ranging from 9,600bps to 10Mbps, are what can be described as current (at the low end) and emerging (at the high end) mass markets for the research data network provider.

The size of this mass market is estimated at around 750,000 individuals, being the population of Australia's research community and of its higher education institutions. We expect that by the end of the decade, every individual in this population will not only have access to RDNs but will be a routine user of at least some of the services/applications listed in Figure 1, in the first four categories. In addition, a substantial corporate and private user community will emerge with direct Internet access.

Figure 1: Research Data Network Market Segmentation

Service Segmentation	Effective End-to-End User Bandwidth Requirement	Size of the Addressable Australian Market*
GROUP 1: GENERAL PURPOSE RDN USERS		
Narrowband Internet services: e-mail; telnet; ftp; news/bulletin board; ASCII-form information retrieval	9600 bps (to any screen with a modem or NTU)	500,000 to 750,000
Isochronous digital audio and video-conferencing, and asynchronous data file transfer, hyperdocument transfer (e.g. Mosaic).	n x 64Kbps (mostly n=2 to 6 i.e. 128Kbps to 384Kbps) (to any late model PC or Mac or PictureTel meeting) Note: X-Windows requires at least 64Kbps access.	200,000+
Distributed collaborative workspace applications — group videophone conferencing with distributed multi-user shared workspace tools, i.e. real time interactive multimedia traffic with low latency	1 to 2Mbps (to meeting room or to a high performance PC)	50,000+
Distributed high resolution 2D imagery, commercial quality video/audio applications incorporating compression (JPEG, MPEG2).	5 to 10Mbps (uncongested 10Base-T Ethernet channels to advanced workstations on the desktop)	5,000 to 10,000
GROUP 2: SPECIFIC PURPOSE RDN USERS		
Time-sensitive visualisation, large file transfer, bursty asynchronous traffic (e.g. GIS, CAD, remote sensing)	10 to 200Mbps	1,000+
Telepresence, low latency on-line access to remote instruments and supercomputers, distributed high resolution 3D and real-time 4D applications	200Mbps to 2 Gbps	200+
Telecommunications and HPC systems engineers who require access to next generation technology	Gigabit testbeds	500+

* Numbers of potential users in research and higher education

The populations set out against applications in the 64Kbps to 10Mbps range are subsets of the 750,000 person market, and reflect temporary constraints based on lack of access, lack of host device capabilities, and/or lack of budget to procure RDN services. Based on interviews we believe these constraints will be largely removed, for services requiring up to 10Mbps bandwidth, by 1997. However it will take several years more for the RDN user population to learn how to make best use of the network facilities available to them.

Current AARNet measures show traffic over the domestic intrastate, interstate and international backbone is now over 800 Gbytes per week.

The international T1 link (1.5Mbps) to the US saturated at busy hour days after it was installed. It is believed that up to half of all AARNet traffic that leaves a region (e.g. the State of Queensland, the ACT) is destined for the Australia-US link. Given that Australian-based researchers say that their research peer group is at least 90% (and often 98%) located outside of this country, the issue of affordable connectivity of adequate (i.e. non-congested) bandwidth to the world is central to Australian RDN planning.

In order to understand the potential for the RDN to be able to accommodate applications requiring peak bandwidth above the present mass market (<10Mbps) applications, we need to attempt to dimension the basic RDN infrastructure during the next few years. But any discussion of Internet/AARNet traffic must be prefaced by several caveats:

- *The number of users is not measurable*

There are no statistically reliable surveys on numbers of users per networked device (host). People guess at 4 or 5 users per device, as an average. A survey by SRI International of the entire Internet population of 726,582 hosts at January 1992, concluded that 98.4% of hosts had only one IP address. However many individuals could share an address in, say, an administrative office LAN environment, or in an undergraduate computer science or engineering laboratory.

We can hypothesise that, over time, most research professionals, academic staff (research, teaching, librarians, specialist support) and higher degree students will be accessing the network from a desktop screen in their personal workspace, so the ratio of users to IP addresses will become 1:1 for this group. For undergraduates, they will have their personal digital assistants by the late '90s, no doubt with embedded telephone numbers. This argues for individual IP addresses with network access becoming mobile rather than fixed (although today's Internet cannot recognise mobile hosts). It may take longer, but this population of users will likely become 1:1 to IP addresses also.

For today's network, perhaps 1.5 users per IP address for the first group and 10:1 for those undergraduates who have been granted AARNet access is a better approximation of reality.

- *The number of hosts is not measurable*

Internet records do not exist for many hosts as institutions have constructed 'firewall technology' at their gateway to the Internet to prevent unauthorised access to their customer premises networks (CPNs). On the other hand, many IP addresses have been issued that may not be associated with active devices. Thus the Internet count of host numbers is a derived function of these two factors, and does not represent a real census of devices.

- *AARNet traffic measurements are not precise*

Some domestic interstate traffic is double or triple counted, but no-one knows what percentage. What is accurately measured is international link traffic in and out. With these qualifications, present measurements suggest traffic entering AARNet (from domestic CPNs and from overseas) is averaging 8 megabytes/week/host, over 100,000 measured hosts in Australia i.e. a total load of 800 gigabytes/week.

Many factors are now influencing traffic growth and network dimensioning projections, including:

- *RDN use by existing users is growing faster than connection of new users*

During 1993, AARNet usage by existing hosts grew faster (81%) than usage due to the rate of connection (58%) of new hosts (see Figure 2 on page 77). This is to be expected, given the youth of AARNet, as the early adopters are still growing in the agility and sophistication with which they ustilise the network's resources. New users do not initially generate high traffic, as they learn to use the basics, such as e-mail and Gopher.

It is key to assess how the number of devices connected to the network will grow, as AARNet's client institutions roll out CPN connectivity. We have interviewed AARNet managers and infrastructure providers at nine research-oriented universities and the CSIRO, to understand the present status of device penetration to their target user population. Broadly speaking, they are half way there now, with target penetration

universally expected to be reached within two to three years.

Thus we might expect the measured number of hosts to at least double by late 1996 to 200,000. This implies a forward growth rate of 25% per annum rather than the 58% experienced in 1993. A less conservative assumption could see 450,000 visible hosts by 1996, which is still lower than recent growth in host connections for both AARNet and the total Internet (see Figure 3 on page 78). However, much of such 'visibility' could be due to existing users of shared addresses, particularly undergraduates, being issued individual IP addresses.

- *Much more bandwidth will be required per user*

Several forces are at work here, both on the demand side and the supply side. New easy-to-use interfaces to the Internet (e.g. Mosaic/World Wide Web) are encouraging both new and experienced users to browse the network. Mosaic users consume 30 to 50 kilobytes of data per page versus 1 or 2 kilobytes for an ASCII e-mail page. Dozens of new Web servers are appearing on the Internet each week.

These new interfaces enable deployment of applications that novice users can adopt rapidly into their workflow design, applications like multimedia and videoconferencing, which will be particularly bandwidth-hungry. The increasing trend toward client/server computing will also have an impact.

AARNet's customers recognise these realities. They all tell us that they are moving to provision their staff and students with access to Ethernet speed LANs and campus backbone networks. This process should be largely complete by 1996. Thus, the user community will be enabled to access all the services set out in the first four segments in Figure 1. Terminal functionality for the large corporate sector will also be deployed in the research and higher education sector, but probably not as quickly (perhaps a five year deployment/scrappage cycle rather than a three year cycle for firms).

Figure 2: AARNet Traffic Growth, 27/9/92 - 19/12/93

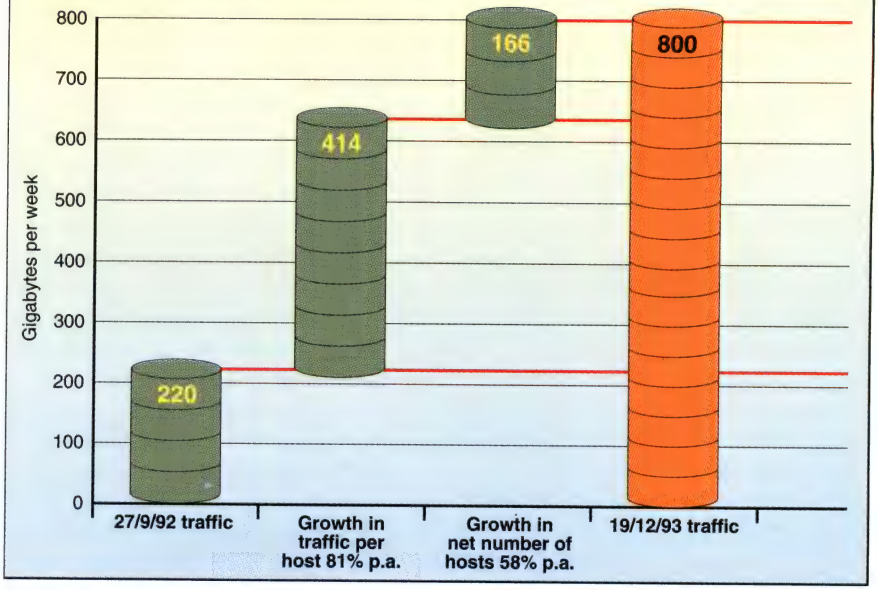

On a more global basis, overseas surveys of large corporate and government users of data networks have revealed a widely-held expectation that the push for higher bandwidth services will cause at least a five-fold increase in intra-CPN and inter-CPN aggregate speeds, between 1993 and 1998.

- *A pause at 10Mbps?*

Current network dimensioning plans, including many university campus networks, do not envisage user access at greater than Ethernet speeds (10Mbps) for at least the next two to three years. This observation is based upon interviews with AARNet management, with infrastructure providers at nine research-oriented universities, and at the CSIRO. The only exception to this is a small community of users, who anticipate gaining 'ATM to the desktop' functionality during 1995/96 — probably at nominal 34-Mbps, 50Mbps, and 155Mbps access rates, but only within the institution's CPN, or at most via an inter-institutional link in a major city.

These users, depicted in Figure 1 as three segments requiring 10Mbps through to gigabit network performance, are estimated to be around 2,000 in number today,

Figure 3: AARNet and Internet Hosts

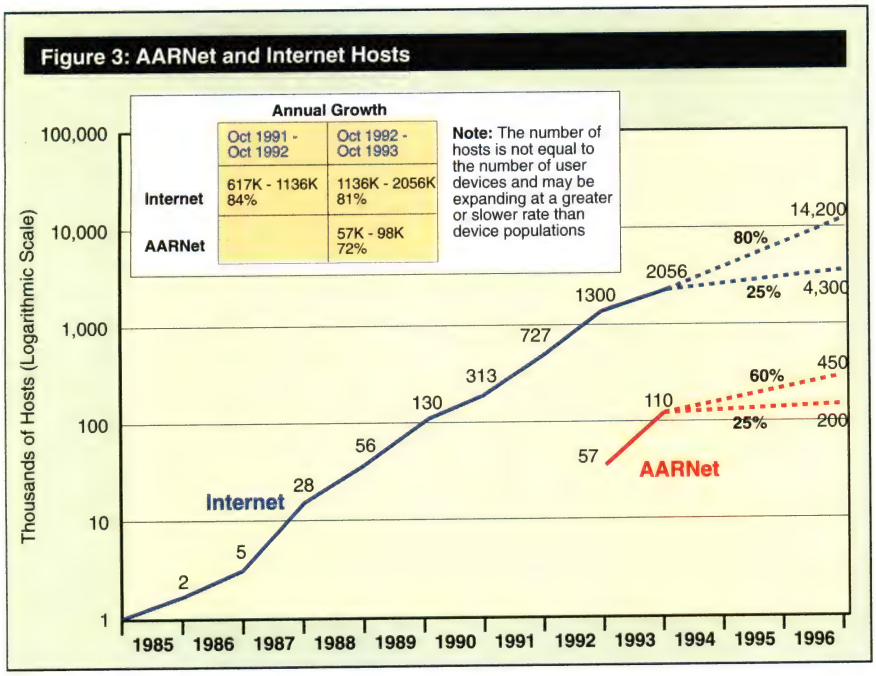

but are expected to increase rapidly over the next 3 to 5 years, particularly due to high performance computing technologies becoming more widely adopted by research and development professionals.

AARNet does not have a brief from its customers/members to dimension its intrastate, interstate or international WAN links to accommodate those requiring a sustained grade of service above 10Mbps. Rather, it has been managed on the basis of continuously dimensioning to support IP traffic at just short of the saturation point i.e. the optimal cost/benefit point for *aggregate* data-flow. Thus its current focus is on how to fund and put in place the massive increase in capacity required over the next three years to handle aggregate traffic of the mass (<10Mbps) marketplace.

Current traffic growth trends suggest a five to six-fold increase in individual user traffic over the next three years, leading to 12 times growth in total traffic on AARNet. This should be thought of as a medium level forecast, but being a minimum scenario for capacity planning purposes. However, a 100-fold increase in total AARNet traffic can be argued as being in prospect (see Figure 4 on page 80), if current cost recovery policies remain in place. That is, cost recovery at the institutional level, rather than per transaction charges 'looking over the shoulder of the researcher.'

Such projections assume a doubling in the number of hosts (and, by implication users), and that individual user traffic could grow 50-fold for this higher demand forecast. The key assumption made here is that the overall increase in application bandwidth could cause an average transaction to require around half a megabit per second of bandwidth (i.e. 512Kbps) rather than the 9,600bps to 64Kbps range for traffic carried by the present research data network in Australia.

- *Implications of pricing policies to keep average annual user costs below $1,000*

If per user volumes were to increase 6-fold or 50-fold, as postulated in Figure 4, what would a volume charging model require to recover costs? AARNet costs in the last quarter of 1993 equated to 30 cents per average transmitted megabyte in the domestic network and 73 cents per megabyte over the international link, both computed at optimal line loads at flat pricing. At the 85:15 domestic:international mix at that time, the average 8 Mbytes/week/host would have been charged 33.5 cents/Mbyte, equating to $140 per annum.

This is a substantially lower price than that being offered by commercial LAN interconnect service providers like Telecom (where even at 50 times the present 8 Mbytes/week volume the Fastpac 2 service averages at $2.37/Mbyte). It is also cheaper than the volume pricing offered by commercial Internet service providers in Australia.

We do not perceive AARNet's proposed per megabyte volume charges, if confirmed in the $0.40 to $1.00 range ($1.50 to US), as being so high as to drastically inhibit per user AARNet demand growth from increasing 6- or 7-fold from current levels i.e. to around $1,000 per user per year. Commercial providers, such as Pegasus Network Communications, are seeing rapid growth of users of narrowband services, with average annual customer usage over $300 today. Beyond the level of $1,000 per year, however, price elasticity effects could become important. Per megabyte charges would need to be around 5 cents to contain average annual user costs to around $1,000.

- *A $200 million a year AARNet by 1997?*

If average user traffic did grow to 427 Mbytes/week, and total AARNet traffic to 85,333 Gbytes/week, could a network be provisioned at 5 cents/Mbyte, at optimal line loads and flat pricing? Derived from the analysis in Figure 4, the implication for network dimensioning is dual 155Mbps links in parts of the interstate backbone and one 155Mbps link to the US.

Our arithmetic suggests that such a backbone plus AARNet overheads could move user charges to less than 5 cents/Mbyte. An overall RDN operating cost of some $200 million per annum would be required, to support some 85.3 terabytes of traffic per week.

This could be feasible if carriers were willing to supply 155Mbps PVC links at, say, no worse than pro-rata pricing to today's 2Mbps links. This implies around $15 million per average domestic intercapital link and $50 million for the two half circuits to the US. A domestic star configuration could consist of single links from Perth/Adelaide, Brisbane and Canberra to Melbourne, and dual links from Sydney. The interstate and international infrastructure might then cost $140 million per annum, leaving $60 million for intrastate tails and network overheads.

We cannot, of course, do more than speculate on whether such a cost structure could be negotiated. However, it would be prudent for RDN planners to commence the process of dialogue with the carriers now, to determine the economic parameters that could be needed within 3 to 4 years.

Several interviewees have made the point that their dataflows are oriented more to Asian and European sources/destinations. A diversity of international links would better preserve Australia's independence of access and pricing policies of other jurisdictions. However, three such full duplex international links (at, say, 50Mbps each) would need to be priced in aggregate at no more than $50 million per annum.

- *A 155Mbps backbone will be technically possible by 1996*

Is it technically realistic to expect carriers to commercially provide backbone links at 155Mbps and higher ATM speeds by 1996? We believe the answer to be in the affirmative, based upon dialogue with Telecom Australia network planners.

Initial commercial 155Mbps Permanent Virtual Circuit (PVC) services should be feasible by mid-1996, given that commercial network systems products, based upon agreed ITU-TS standards, are becoming available to carriers in 1994. Telecom's Experimental Broadband Network (EBN) is expected to deploy 155Mbps PVC infrastructure in the first quarter of 1995, which should support a limited number of 155-Mbps service interfaces later in that year. Switched Virtual Circuit (SVC) standards, systems and service availability are expected around one year later. From the supp-

ly side, then, it appears that carriers could provide high bandwidth links to support very high aggregate RDN traffic growth over the next three years and beyond.

The HPCC Community

ASTEC asked Cutler & Company to pay special attention to the needs of the high performance computer and communications (HPCC) community. We characterised such users as those requiring 10Mbps or higher bandwidth, on a reliable basis, for sustained traffic and for peak rate bursty traffic.

The task is to build a case for the necessary and sufficient RDN functionality and capacity such that bandwidth in the local, regional, and wide area network is not a bottleneck or constraint in the particular application or research system design. Further, such bandwidth must be available at a price that is reasonably reflective of underlying cost of provision of the network (as is AARNet's current philosophy), rather than at high premiums that might be adopted by carriers to manage the rate at which their customers can economically justify switching from current lower bandwidth data services (that are usually highly profitable, e.g. the ISDN service).

In the early years of so-called 'gigabit networking,' costs will be high, particularly for private gigabit networks compared to public gigabit networks. In essence, AARNet operates today on the public network model in that it supports a broad range of applications, is widely deployed, and has many users. It is likely to deploy gigabit network capacity during the next few years on the public gigabit network model, to handle aggregate traffic.

Telecom's EBN will be deployed on the private gigabit network model, in that it is expected that it will be dimensioned to provide a relatively small group of application developers and network engineers with adequate, reservable end-to-end bandwidth to suit their research design.

HPCC Applications

For an application to require data rates in the 100Mbps to gigabit per second range, it must:

- Be geographically distributed;
- Require transmission of large amounts of data; and
- Require low latency.

Even where large (multi-gigabyte) files are to be transferred, if transfer time is not critical then lower speed networks will suffice. Even for small data sets, end-to-end latency is determined by speed-of-light delays. Latency also impacts on choices of transmission paths. For this reason, intercontinental satellite paths from Australia to, for example, the US have been seen to pose problems for gigabit networking.

One analysis suggests that a 155Mbps stream over satellite would sustain 670 microsecond delays, compared to 150ms across the submarine fibre cable. Even for fibre, the engineering task to set up 20 Mbyte buffers is still to be done. Australia's latency issue is more acute than for, say, New York to San Francisco traffic, where 50ms two-way propagation delays are now being addressed by network engineers working in the various US gigabit testbed programs.

It is the combination of large data transfers and low latency that will demand gigabit networks for individual applications. Our interviews within the Australian research community, and reading of the literature on work-in-progress at the gigabit network test beds in the US and the EU, suggests that applications likely to require communication speeds in the hundreds of megabits and gigabits per second range are fundamentally linked to systems in which high performance computing is required (see 'NREN — Setting the Network Clock to Fast Forward' in the November 1992 edition). Such applications include:

- Distributed supercomputing;
- Remote visualisation in near real-time; and
- Telepresence (virtual reality).

Many applications being discussed today within the Australian research community are those involving high definition, three-dimensional imaging, where a user needs to interact with the image in real time. The user may want to rotate the image on the screen, or vary a parameter of the underlying model or simulation whilst observing the effect of such a change. Such human-machine interaction will require the transmission of 15 to 30 frames (screens) per second with sub-one second latencies. Gigabit networks are required when these images and animations must be generated remotely from the computational source.

All potential applications requiring gigabit networks are still in the laboratory at this point, so the economic benefits of operational systems have yet to be demonstrated in commercial performance terms. However, several examples of present experimental systems can be described to illustrate the network performance requirement. They include teleradiology and remote consultation, dynamic radiation therapy planning, complex neurosurgery, and interactive 3D rendering of multiple, very large earth science data sets. There is significant Australian-based research interest in these new applications.

It is worthwhile to note that we have found many research domains to be less

sensitive to time delays in network delivery of data than might be at first supposed. Most, when pressed, could happily live with a reliable, uncongested 10Mbps service, that was widely available. Examples include biological molecular modellers and spatial information services.

Regional Connections

Many research problems in computational science and engineering are driven by access to very high performance computers. The scale of HPC machines and the pattern of their deployment in Australia is likely to be the critical influence on the requirement for very high bandwidth RDN infrastructure in Australia during the next four to five years.

High performance computing is currently the preserve of academic and government research, based on a series of regional centres. Supercomputing is in its infancy with Australian industry. Most industrial R&D and operational users are at the PC and workstation level, BHP has an entry level Cray EL. Larger research models are run for industry and government clients on university or CSIRO production machines (such as the Cray Y/MP's at CSIRO Melbourne and the University of Queensland, and at the RMIT/University of Melbourne Ormond facility). The Bureau of Meteorology and Department of Defence have in-house supercomputing facilities.

Industry outreach programs are in a start up phase, such as at VisLab at the University of Sydney, at Ormond in Melbourne, at QSL/CSIRO in Brisbane, at ANSAMS in Sydney and at ACCI in Melbourne. Clients come to the regional HPC centre rather than seek remote access via networks at this point in the experience curve.

Federal Government funding has supported regional academic centres for supercomputing that have emerged in recent years and State Governments are supporting outreach initiatives based on their regional centres and naturally enough see merit in encouraging a local training and industry support capability. Their emphasis on network access to HPC facilities is consistent with that of the local universities that host these facilities: build up local and regional access, at relatively low cost over high capacity microwave or short fibre links.

Thus we observe activity in many cities to set up metropolitan area networks that, inter-alia, will link HPC services. Brisbane has QUESTNet, Melbourne has VicNet, Adelaide has SARDNet plans, and Perth is seeking funding for PARNet. Federal Government funding for the Research Data Network CRC is being keenly eyed as a one-time fillip to regional infrastructure building: both the $3.4 million allocated to AARNet upgrades and the $3.0 million to be used to fund local access to Telecom's EBN.

It is a point of speculation as to how the regional HPC communities would seek to use the EBN's 155Mbps connectivity. We see some interest in using the regional ATM switched EBN infrastructure to gain 'metacentre' benefits within the region. It is less obvious as to how much HPCC-related traffic would flow between, for example, Brisbane and Sydney, or Melbourne and Canberra. Only one research project, to our knowledge, has been proposed that would apply high data loads on intercapital EBN links, being the RDN-CRC's Sydney-Canberra-Adelaide proposal, linking several Connection Machines in a computer science research setting (rather than a production environment).

There is obvious rivalry between the various regional academic groupings at this time, which should not be dismissed by those seeking to plot a course for rapid diffusion of HPC-based technologies and capabilities into Australian research and commerce. The impetus is towards a policy similar to that adopted by many States in the US: medium sized centres in regions where users can physically visit as well as affordably access electronically. This is the Alabama or North Carolina model, rather than the NSF/Federal Government-supported National Supercomputing Centres and the National Laboratories where state-of-the-art supercomputers are deployed.

Thus one can foresee a steady building up of local capabilities in the (at least) five regional centres. A variety of medium-powered HPC machines (both scalar/vector and parallel architectures) connected by very high bandwidth private data networks would be in place in each region. In this way, five or six vector machines of 5 G-FLOP capacity and five or six parallel machines of 50 GFLOP peak capacity might be deployed in Australia during the next three years or so, serving hundreds of research institutions and industrial clients as the penetration of HPC functionality into the wider community gets under way.

Central

If the above scenario plays out, what will be the requirement of the HPC users for high bandwidth long-haul interstate and international links? Certainly less than would be required for the alternative, a centralised National Supercomputing Centre with only modest experimental machines in the regions, but not zero. There will still be classes

Figure 4: RDN Traffic Projections

	AARNet (End 1993)	Medium Case (End 1996)	Higher Case (End 1996)
Total hosts	100,000	200,000	200,000
Traffic/host	8 Mbytes/week	50 Mbytes/week at 80-85% per annum i.e. the present rate of growth	427 Mbytes/week/host — by calculation
Average bandwidth	Assume 9,600bps	60Kbps — by calculation, if connect time/day remains constant	Assume 512kbps
Average connect time/day (5 day week)	22 minutes	Assume 22 Minutes	Assume 22 Minutes
Total traffic	800 Gbytes/week	10,000 Gbytes/week 132% per annum growth (i.e. the present rate of growth continues)	85,833 Gbytes/week 375% per annum growth
AARNet Backbone Capacity	2Mbps (saturating)	25Mbps required — implies 34Mbps MAN/WAN, international link	213Mbps required — implies 2 x 155Mbps ATM links required, for main domestic trunks, and one 155Mbps international link

Notes and Assumptions

1. AARNet estimates of end '93 traffic/host/week are accurate to ±20%, at 8 Mbytes. Average user access to the RDN is assumed to be 9600 bps. Thus 22 minutes/day connect time is a derived value. It is not experimentally validated.
2. Both forecast scenarios assume 22 minutes/day remains a reasonable representation of user work patterns, in a communications environment that is swiftly moving from text-based to audio-visual and hypermedia-based.
3. The medium case extends CY1993 growth/host at the same rate to end CY1996. It implies near universal access to AARNet being available at ISDN bandwidths (n x 64Kbps), for the mass research and academic user population. Perhaps only undergraduates would be accessing at lower (modem-based) speeds.
4. The higher case assumes average bandwidth/transaction will be 512Kbps. This implies widespread audio-visual traffic and high resolution visualisation traffic. The question is not whether this state will be reached, but rather it is when — it could be a year or two after 1996.

of problems that will need access to the greatest CPU and I/O and mass storage performance available in the world. The public policy issue for Australia is whether a facility with such capabilities, costing, say, $30 million per year to build and run, can be justified. It would be a National Facility, comparable to the Australia Telescope or the (as yet unfunded) HIFAR nuclear science facility.

Our impression is that the constituency for such a Facility has yet to make its case. Until it is made, then an overlay private very high bandwidth network between such a Facility and the existing regional HPC centres cannot be contemplated. Rather, the inter-regional HPC traffic would be carried by the mass market RDN (AARNet), which will be likely to have the nominal bandwidth required for HPCC traffic in its interstate backbone, but which will present the HPC user with congestion and severe system performance degradation, from time-to-time.

Similar comments apply to HPCC traffic to/from overseas centres. The Internet may well allow international sharing of 'spare cycles' but it will not offer any performance guarantees for such distributed supercomputing systems until such time as reserved or priority allocations of bandwidth can be offered by public gigabit networks.

We understand considerable effort is being devoted by the ATM developers and standards-setters worldwide to defining and providing such service. This is driven by grade-of-service requirements of the <10-Mbps interactive multimedia users (an emerging mass market for the late '90s), rather than by the far less numerous HPCC users. However public ATM networks may provide the answer for HPCC system performance within a five to six year outlook. Meantime, Australian domestic HPCC users can adopt Telecom's EBN — provided they are co-located with an EBN tail, in the right place in five Eastern Australian cities.

A final note of caution should be sounded though for the regional HPC centre approach described above. It is indeed a user/application driven model, that has been seen in the US to be an essential part of diffusion of HPC into industry nationwide. However the process in the US started with the NSF National Supercomputing Centres initiative. These centres provided testbeds for development of new computing and communications systems that have turned into commercial products, both hardware and software.

The question for Australian product developers, particularly of software utilities and tools and applications, is whether they can access global markets on a timely basis if they do not have access to very high performance computing facilities and to network testbeds which are operating at gigabit speeds. The Australian approach seems to be a more modest 'fast follower' idea i.e. 5 GFLOP rather than 100 GFLOP computers and a 155Mbps EBN rather than the 622Mbps ATM/800Mbps HIPPI combination now being used in the US.

This time-lag or 'late to market' factor for gigabit networking seems to be two to three years. For example, Telecom's EBN will be deploying 622Mbps technology some three years later than the US gigabit testbeds (see Figure 5 on page 82). The combination of regionalised HPC centres and lack of competitive and client pressures on local carriers to invest in gigabit network R&D, would seem to argue for the lag increasing rather than the gap being closed.

The above factors will have significant implications for network capacity planners in the late 90s, to the extent that deployment of operational HPCC-based systems, that could serve overseas markets from an Australian base, may not be able to capture 'first mover' market advantage. This is of course not the principal concern of RDN providers in a direct sense, but may be argued to impact longer term on the competitiveness of Australia as an R&D base, and hence the scale of activities possible for RDN clients. An indicator of the potential loss to the nation from not pursuing both regional HPC centre and leading-edge HPCC research strategies can be seen in the following communication we have had with the leader of the gigabit testbed project at the National Centre for Supercomputing Applications at Champaign-Urbana, Illinois:

'What we are trying to do is to make sure WAN speeds match well with LAN speeds i.e. so that the only difference to the user between using a resource at one location vs another (or a combination of two separate remote resources) is speed of light. What is stopping us right now is mainly host I/O i.e. machines that can really drive HIPPI (800-Mbps). Cray can do it now, SGI and CM.5 should do well by summer with the release of new HIPPI interfaces. Other workstations (Sun, DEC, HP, etc.) will probably not be far behind SGI. We have dozens of applications that require tens of Mbps, about a dozen that require 200-800Mbps, and a handful that require several Gbps.'

Demand in Australia for RDNs that can supply gigabit bandwidth to individual transactions will depend upon 'dozens of applications' being adopted or created by local researchers and end users. There are HPC applications currently in use or in development in Australia and there would be merit in ASTEC or another agency addressing issues for a national high performance

Figure 5: Australian Bandwidth Availability Projections

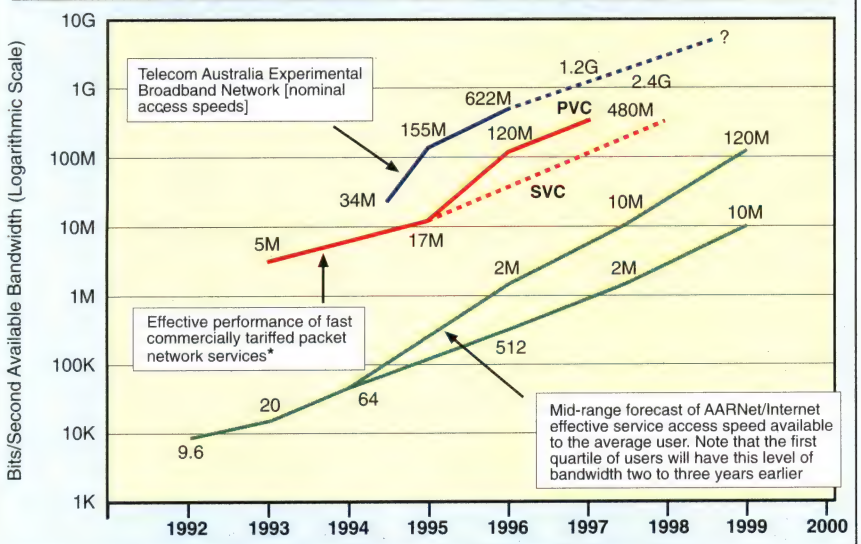

*Fastpac 10 and 34 access switched service to 1996 and beyond; ATM-based Permanent Virtual Circuit service from mid-1996; ATM-based Switched Virtual Circuit service from mid-1997

computing strategy, which could then provide a more relevant context in which to plan and to construct very high performance network infrastructure.

AARNet's Challenge

The Australian RDN of tomorrow should be put in place to support both the aggregate demand load and the peak bandwidth required by services aimed at a mass market in the research and higher education sectors. This market is approximately three quarters of a million people and the effective bandwidth required by, say, 99.5% of current users, is accommodated within clear (uncongested) Ethernet channel speeds i.e. up to 10Mbps.

This will prevail for at least the next three years, as a substantial education and training effort will need to be undertaken to enable the majority of academic staff and postgraduate students in higher education institutions, and all staff in research establishments (both government and private sector), to be adequately empowered to use the functionality being provided to them. People who today view their screens as standalone word processing, spreadsheet and database machines will need much assistance to adapt to the network culture of client/server computing, e-mail worldwide, global information search and retrieval, and the use of computer-supported collaborative work (CSCW) technologies.

It is likely that aggregate capacity will be provisioned by AARNet in the major cities and in the interstate and international links at bandwidths sufficient to provide non-busy hour opportunity for individual applications requiring 10Mbps speeds. This is highly likely in the 10 to 34Mbps Fastpac range and probable in the 34 to 155Mbps ATM range during the next two to three years.

However, by three years from now, substantial numbers of researchers are likely to be feeling the limits of 10Mbps networks, as do some of the members of the HPCC community today. This small 'gigabit networking' community (less than 0.5% of the user population) will not be optimally served as their needs for bandwidth are for much higher (10 to 200 times) speeds than those needed by the mass market. We cannot be sure how many users, requiring 20 to 120-Mbps performance, will move into this community by 1997, but there likely will be tens of thousands of researchers with this requirement by then.

We do not foresee any technical problems in accommodating demand for higher bandwidth applications, as the first commercial gigabit wide area networks switching traffic at nominal 155Mbps speeds should be deployed in and between all major Australian metropolitan areas by that time. However, the tariffing of such services by the carriers is of critical concern (there is no switched service tariffed today in Australia at greater than 10Mbps access speeds, whereas 34Mbps and 45Mbps domestic switched services can be purchased now in the US and the UK).

It is critical for Australian RDN planners to recognise that demand-driven aggregate traffic growth will drive AARNet into gigabit networking, on a public carrier model. Demand may require a $200 million per year AARNet by 1997, rather than the $15 to $20 million operation being contemplated today, if commercial prices are to be paid to carriers for WAN links.

Providers of network capacity will always dimension for aggregate traffic at busy hours rather than the peak bandwidth requirements of an individual HPCC user. Such an HPCC user cannot be given the guarantee of a minimum acceptable bandwidth for the duration of his or her network transaction that is needed to be more sure of achieving desired end-to-end HPCC system performance. So even a geographically-restricted (i.e. major city) availability of high bandwidth access to the RDN backbone will not necessarily solve HPCC user needs. Telecom Australia's Experimental Broadband Network may become the only platform in the near term in Australia that will combine the low cost of access and usage that one might expect from the RDN dimensioned for the mass market with the ability to reserve guaranteed bandwidth for the duration of an HPCC transaction.

Up to 30% of all AARNet traffic is international. The implications for trans-Pacific capacity in a few years from now are for demand at as much as 100 times present constrained volumes. No carrier has tariffed trans-Pacific T3 (45Mbps) links, let alone the 155Mbps link that could be required. Whereas it is a reasonable proposition that the RDN should be user-funded, this may not be achievable for the international link at the earliest feasible date (i.e. 1995) without public policy action. Such action might seek to encourage the Australian and US-based carriers who would have to be involved (providing a half-circuit each) to provide service on an interruptible, marginal cost of provision basis, using excess fibre capacity that exists today, until such time as that capacity is required for other carrier services. By that time, around 1997, the carriers may have provisioned more international cable (low latency) capacity, and should have filed tariffs for such high bandwidth services.

RDN capacity planning is totally a function of aggregate demand from large numbers of geographically dispersed research and higher education users. Both AARNet and State-based regional RDN consortia are engaged in provision of infrastructure to service this mass market. To date they have coordinated their funding and have developed an integrated single network. However, there is no certainty that this coincidence of interest will prevail should Telecom or Optus or other value added network providers seek to deal with one or more of the regional groupings. AARNet's scale economies, in aggregating demand to justify higher bandwidth trunks that deliver lower per megabyte costs, could be threatened by such decentralisation of capacity buying decisions.

The major challenge for AARNet will be to be financially credible to the carriers, as it signs contracts for up to three years committed purchase of tens and hundreds of megabits/second of trunk and tail loop capacity.

Cutler & Company is a Melbourne-based consultancy specialising in communications.

The ACE ROUTER is the World's 1st fully Integrated routing solution for LAN Traffic over ISDN.

COMPLETELY INTEGRATED

Australian Made and Supported

(No external Terminal Adaptor required)

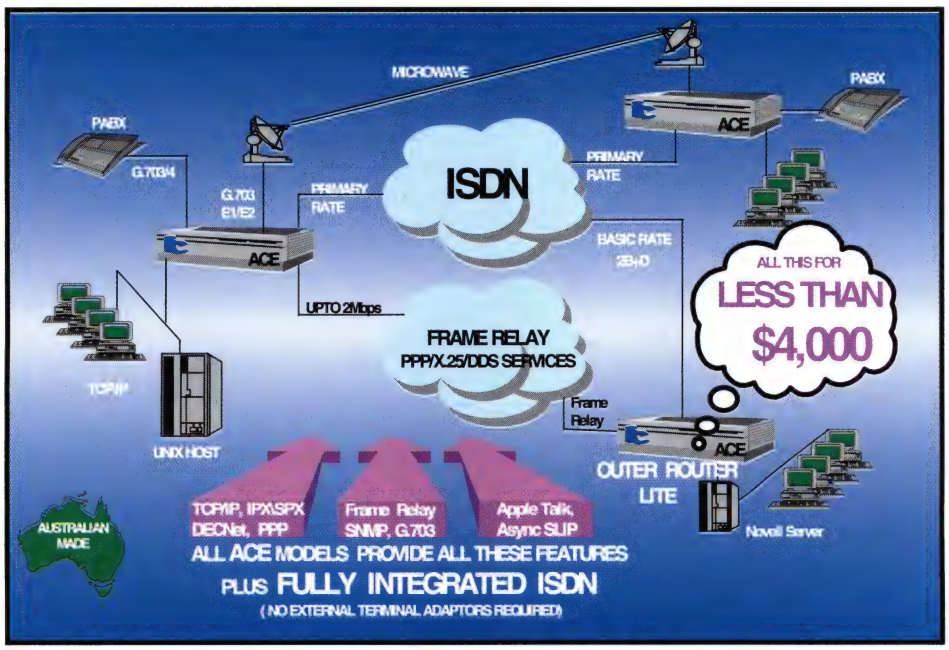

Interlink Communications

Head Office
PH: 02 975 2577
FAX: 02 452 5397

Melbourne Office
PH: 03 822 9831
FAX: 03 822 9980

Canberra Office
PH: 06 257 1885
FAX: 06 257 1998

BREAKING THE LAN SPEED RECORD

OLICOM POWERMACH TOKEN-RING ADAPTER DRIVERS

When we introduced our redesigned adapter drivers last year, they were the fastest in the world. Now, with PowerMACH, we've made them 30% faster! No wonder they're still running rings around the competition. Just look:

At small frame sizes such as 512 bytes—which account for over half of all Token-Ring traffic—each time the market leader sends five frames, PowerMACH sends nine!* What's more, PowerMACH optimizes CPU utilization so your computer has more time to perform your applications, and spends less time playing mailman.

And when it comes to flexibility and convenience, PowerMACH is second to none. For instance, PowerMACH features a comprehensive SNMP desktop agent and is prepared for DMI—the new industry standard for desktop management. There's also a special Windows utility that provides your network manager with important adapter information.

In short, high-performance hardware needs drivers and features to match. And with Olicom...well...boot up and watch the competition take a hike!

* Novell Perform 3 Benchmark test using latest version ISA adapters and drivers.

You can download the latest version PowerMACH drivers absolutely free simply by calling one of the following services:

Olicom Bulletin Board Services
USA: +1 214 422 9835
Europe: +45 27 01 00

CompuServe®
(type "GO OLICOM" after login)
USA: +1 614 547 0802
Europe: +44 272 255 111

olicom

Australian Distributor for Olicom: FORCE TECHNOLOGY PTY LTD, P.O. Box 1516 • Dee Why, NSW 2099 • Ph: (02) 971 1000 • Fax: (02) 981 1932

Building a Better Token Ring Network

MAUs, hubs, switches, 100Mbps technologies; the new products coming onto the market add up to some real choices for builders and managers of enterprise Token Ring networks.

Form ever follows function. Louis Sullivan may have had the early high-rises — not corporate networks — in mind when he penned that simple imperative nearly a century ago, but the architect's directive holds just as much truth for today's Token Ring planners as they come to terms with a growing array of equipment and design options.

The arrival of every new and improved technology — first modular Token Ring hubs, now LAN switches, and soon high-speed (100Mbps) hubs and adaptors — increases the temptation to replace the multistation access units (MAUs) that served as the first foundations for Token Ring LANs. To hear vendors tell it, that's exactly what network designers are doing.

A quick glance at the typical corporate network landscape reveals a slightly different reality. To be sure, modular hubs with high-speed backplanes are supplanting bare-bones MAUs at strategic points in many networks. Token Ring switches and 100Mbps products also will be deployed at key pressure points. But in the less hectic parts of a network — such as remote branch offices with limited numbers of users — MAUs remain the rule rather than the exception.

The main design challenge for builders of corporate-wide Token Ring networks is to draw up a blueprint that accommodates current and anticipated user needs as efficiently and inexpensively as possible. More often than not, designers are likely to find that the best solution lies in a mix of old and new technologies.

Network planners will have the usual metrics to help make their decisions — comparative costs per port, performance and application requirements, cabling limitations, and the like. But to make things a little more interesting, Token Ring deployment involves a few more issues that need to be taken into account. One concerns network management: designers must decide not only which management approach to take (IBM's LAN Network Manager, which is tailor-made for Token Ring, or the more general SNMP platforms) but also where management functions should be deployed in the network. There's also the question of whether to use active (as opposed to passive) components to negate the effects of signal attenuation, or jitter. And then there's the problem of how to deal with beaconing, which remains a prime cause of outages on Token Ring LANs.

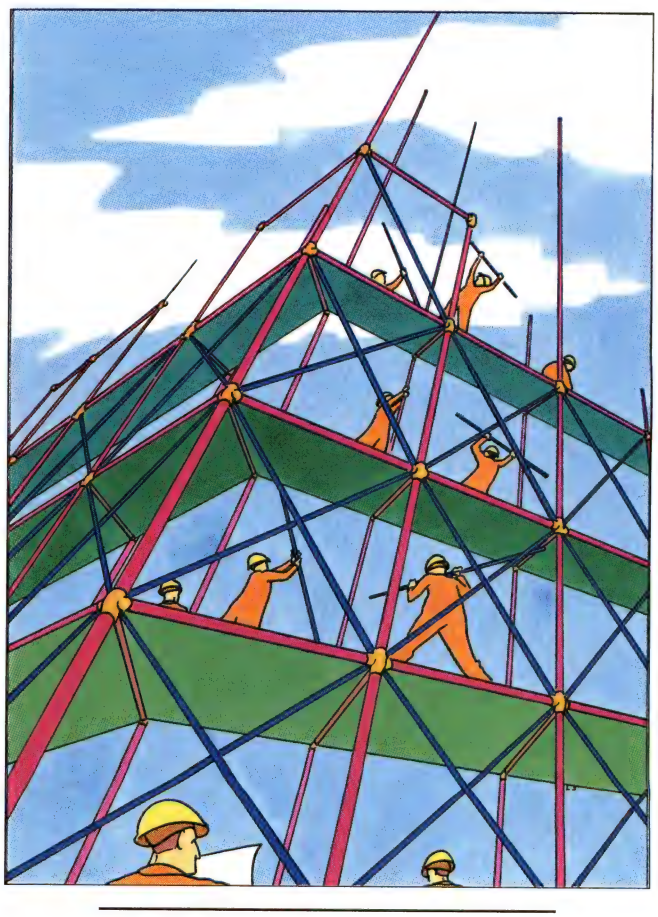

Freedom from Choice

When the IEEE 802.5 standard defining Token Ring debuted back in 1985, net managers didn't have to spend much time thinking about design options. Back then, the only way to build Token Ring LANs was to deploy 8228 MAUs from IBM. MAUs are basic wiring concentrators that typically come with eight ports for PC connections and a ring-in/ring-out port to link to other MAUs to form a larger LAN segment.

The 802.5 standard allows each Token Ring segment to accommodate a maximum of 260 end-stations. (Cabling distance limitations tend to restrict the number of end-stations even further.) Larger networks require the installation of multiple Token Ring segments that are then connected via two-port bridges, which along

300% faster than Cisco.
50% faster than Olicom.
60% faster than IBM.
50% of the Cost.

Introducing the ACSYS TokenSwitch 3000 Bridge/Router

"It's the fastest Token Ring bridge we've ever tested." *Kevin Tolly, The Tolly Group*

The numbers are in. The highest performance Token Ring bridge/router is the new TokenSwitch 3000 from ACSYS. It's the only bridge that is wire speed at 64 byte packets and up – at a per packet cost that's half the competition's. Highest throughput rates and ultra-low latency make it the ideal choice for client/server applications requiring fast response time...and for reducing dropped sessions if you have integrated SNA onto the LAN. TokenSwitch 3000 is IBM® compatible, supports both LAN Network Manager™ and SNMP management. It combines the sophistication of IPX™ and TCP/IP routing with the speed and simplicity of source/route bridging.

Get your free copy of The Tolly Group benchmark and our free **Token Ring Performance** white paper.

Call ACSYS today. You'll be 100% impressed.

Any Questions?
Call 1-800-GO-ACSYS

Singapore
Ph Int 65 432 0311
Fx Int 65 480 4535

Australia
Ph Int 61 2 299 2799
Fx Int 61 2 299 2807
Toll Free 008 63 3337

New Zealand
Ph Int 64 9 634 4580
Fx Int 64 9 634 4586

ACSYS

An AMP Company

Figure 1: Backbone in a Box

Modular Token Ring hubs feature a collapsed backbone architecture that uses a high-speed bus to forward traffic between plug-in modules. Each module can be set up as a self contained LAN segment (as with Networks #1 and #2 below) connected by bridge modules, or multiple modules can form a larger segment (as with Network #3).

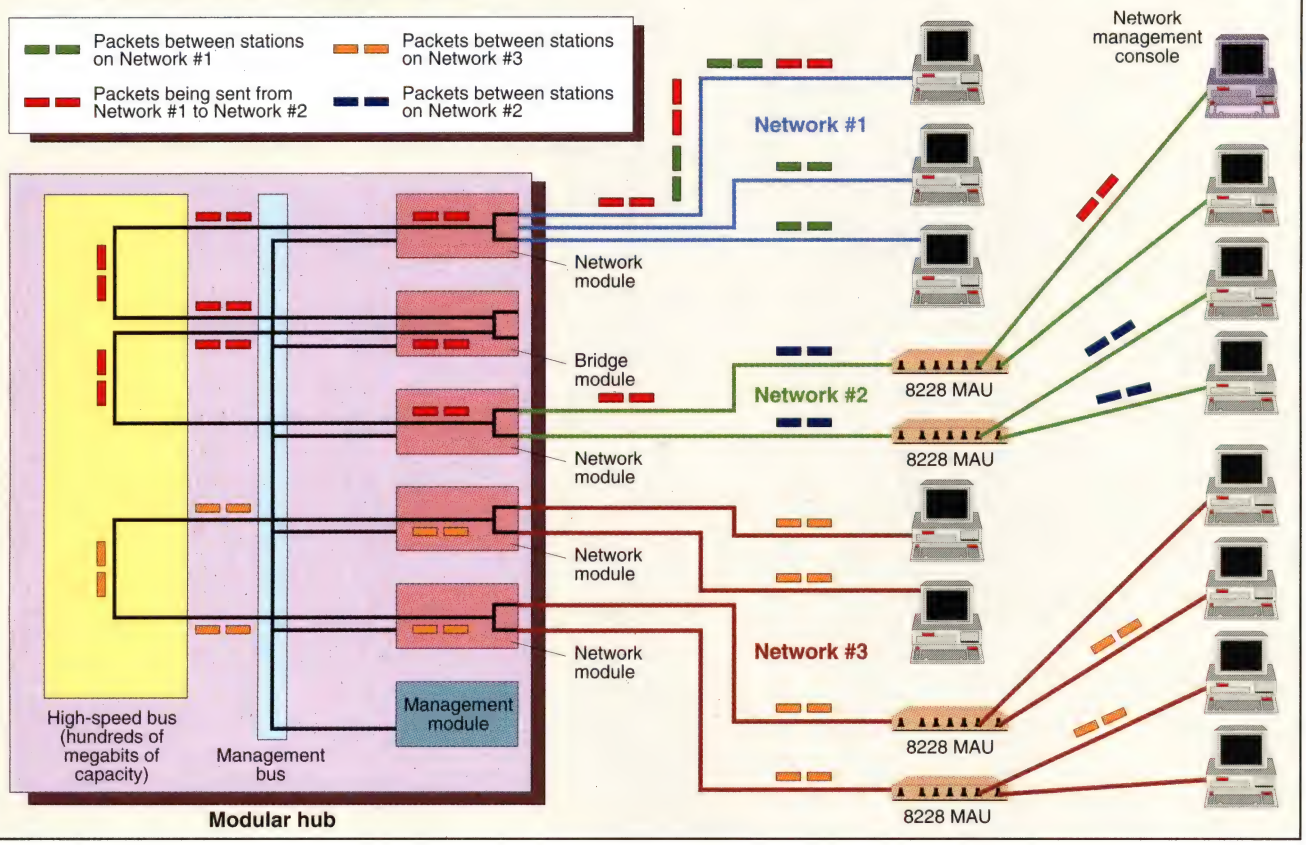

with MAUs served as the basic building blocks for early Token Ring networks.

MAUs like the IBM 8228 are dumb devices, which means they cannot handle functions like network monitoring or management. For this reason, Token Ring bridges typically are equipped with management agents that gather network data for use by a management console running IBM's LAN Network Manager software.

The simple MAU and bridge design worked well enough for first-generation Token Ring networks. And in 1990, IBM improved on its original design from a management standpoint with its release of the 8230 CAU (Controlled Access Unit) and LAM (Lobe Attachment Module) products. CAUs include network management agents, enabling them to perform monitoring functions previously handled only by bridges. With the CAU/LAM approach, designers can build a manageable Token Ring LAN by connecting multiple passive LAN concentrators to a single master CAU. Several vendors now offer variations on this 'smart MAU' approach (see Table 1 on page 88).

One basic problem with Token Ring's first-generation design approach is its limited ability to scale up to handle larger networks or higher bandwidth demands. The distributed, bridged architecture holds up for low-load applications, such as electronic mail and file and print services. But with the advent of more powerful applications, such as those based on the client-server model, net managers need a way to increase the amount of bandwidth available to each user. That means breaking networks into ever-smaller bridged segments.

In networks with multiple bridges, network managers typically deploy a two-tiered architecture in which bridges are connected to a 16Mbps ring that serves as a backbone. The two-tier design limits network delay by guaranteeing that packets have to cross only two bridges to reach their destination — one from their own ring onto the backbone, and the other from the backbone to the destination ring.

The two-tier approach undoubtedly boosts performance for MAU-based Token Ring networks — which is why it's the most common topology found in large Token Ring sites today. But it still presents problems that become more acute as networks grow. First, two-tier networks are complicated to manage, because management information must be collated from multiple distributed segments. Second, reliability can be an issue, particularly when PC-based bridges are deployed. Adding bridges also increases the per-port cost of Token Ring LANs, but the biggest limitation of the two-tiered architecture is on the amount of performance it can deliver. If too many rings are connected to the 16Mbps backbone, the backbone ring can become a bottleneck. Congestion is increased when companies run applications enterprise-wide, as more are now doing — a 16Mbps backbone just can't cope with that much traffic.

Paradigm Shift

The need for more bandwidth is the primary force driving demand for modular Token Ring hubs, chassis-based products holding slot-in hardware modules that supply network connections, bridges, or management facilities.

Modular hubs use the same basic design principles as the two-tiered bridged architecture, in that LAN segments contained on one or more 16Mbps network modules are connected via bridges and a backbone. The key difference is that most modular hubs feature an internal high-speed bus that serves as a collapsed backbone offering hundreds of megabits per second of bandwidth (see Figure 1). There are exceptions to this rule. Some hubs, like the one from and RAD Data Communications, use a passive backplane design based on a single 16Mbps internal bus.

Table 1: A Selection of Smart MAUs

VENDOR	PRODUCT	NODE PORTS/ RIRO PORTS	STACKABLE MANAGEMENT MODULE	SNMP
Bytex Techway (02) 975 2122	7715 TR	12, 24 or 36 RJ-45/2 DB-9	None	Yes
Cabletron Systems (02) 950 5900	TRXI series	12 or 24 RJ-45/2 ST fibre, RJ-45 or DB-9	None	Yes; RMON
Crosscomm Scitec (02) 428 9555	XL-TRSM, IBN-TRSM	16 RJ-45/None	None	Yes: RMON
Digital Equipment Corp. (02) 561 5252	DECMAU 900TL	8RJ-45/2RJ-45	None	Yes
ADC Fibermux (02) 975 1499	SnapLAN-FX520T	5 RJ-45/None	Manages up to 3 dumb MAUs; comes with 12 RJ-45 ports	Yes
Fibronics International Olivetti (02) 368 1988	FR9232/16, FR9232/32	16 or 32 RJ-45/None	Manages up to 16 dumb MAUs; comes with 12 RJ-45 or 24 RJ-45 or RJ-21 ports	Yes
IBM 13 24 26	8230-001/002 Controlled Access Unit	None/2 IBM-type	Manages 4 dumb MAUs; comes with 20 RJ-45 or IBM-type ports	Planned
Madge Networks DDP (03) 694 6711	Smart CAU Plus	4 DB-9/2 IBM-type	Manages 8 dumb MAUs; comes with 20 DB-9 or RJ-45 ports	Yes; RMON
	Smart 16/4 AT Hubcard	10 RJ-45/None	None	Yes
Olicom Force Technology (02) 971 1000	Controlled Access Unit OC-3610	20 IBM-type or RJ-45/2 IBM-type	Manages 3 dumb MAUs; comes with 20 RJ-45 or IBM-type ports	Yes
	Controlled Attachment Module OC-3617	10 RJ-45/None	Manages 2 dumb MAUs; comes with 20 RJ-45 or IBM-type ports	Yes
Proteon (02) 955 8555	Series 70	8 IBM-type/None	None	Yes
Racal-Datacom (02) 936 7023	Internext INXlink-CAU	10 or 26 RJ-45/2RJ-45	None	Yes
RAD Data Communications Dataplex (03) 210 3333	S-TAU	8 DB-9 or RJ-45/None	None	Optional
	S-TAU/RJ	8 RJ-45/None	None	Optional
	X-TAU	8 RG-59 or RG-62/2 RJ-45 or IBM-type	None	Optional
	F-TAU	8 SMA or ST fibre/2 RJ-45 or IBM type	None	Optional
Standard Microsystems Corp. (02) 238 2206	Elite Series 4000C	8 IBM-type or 16 RJ-45/None	Manages 32 dumb MAUs; comes with 8 IBM-type or 16 RJ-45 ports	Optional
SynOptics Communications (03) 853 0799	2712B/2712B-F passive workgroup cluster	16 RJ-45/2 IBM-type or ST fibre	Manages 4 dumb MAUs; comes with 16 RJ-45 ports	Yes
	2715B/2715B-F active retiming workgroup cluster	16 RJ-45/2IBM-type or ST fibre	Manages 4 dumb MAUs; comes with 16 RJ-45 ports	Yes
3Com (02) 959 3020	LinkBuilder TR Series Stackable Token Ring Hub	16 RJ-45/None	Manages 15 dumb MAUs; comes with 16 RJ-45 ports	Yes; RMON
UB Networks (03) 693 8200	Access/Trax	22 RJ-45/2 RJ-45	Manages 5 dumb MAUs; comes with 24 RJ-45 ports	Yes

* For 16Mbps networks; distances in metres except where noted

Along with improved backbone performance, modular hubs offer several other benefits. For instance, hubs are available with or without management modules; such modules add several thousand dollars to the hub's price. Management facilities allow networks to be configured with more flexibility than MAUs. Network managers can link modules into larger segments or split them into smaller ones simply by issuing commands from a network management console.

The hub's high-speed bus can be used to connect either internal (slot-in) or external (standalone) bridges linking different LAN segments. Bridges can be linked to MAU-based segments, giving users a way to preserve investments in older gear. Cabletron

LAN NETWORK MANAGER	MAXIMUM LOBE LENGTH CATEGORY 3UTP/ CATEGORY 5 UTP/STP*	MAXIMUM NODES CATEGORY 3 UTP/ CATEGORY 5 UTP/STP	BEACON RECOVERY
No	100/100/100	72/72/260	No
No	100/130/200	100/100/250	Yes
Yes	132/250/not supported	150/300/0	Yes
No	44/65/100	256/256/256	No
No	100/100/150	72/72/260	Yes
Yes	40/100/150	260/260/260	Yes
Yes	100/100/100	72/72/260	No
Yes	Not supported/100/145	132/132/260	Yes
No	75/100/not supported	250/250/250	Yes
Yes	100/100/145	72/132/260	Yes
Yes	100/100/145	50/50/50	Yes
No	100/100/145	50/50/50	Yes
No	60/100/145	250/250/250	Yes
Optional	253 metres/253metres/253 metres	260/260/260	Optional
Optional	84 metres/141 metres/not supported	72/140/0	Optional
Optional	100 metres/100metres/100 metres	30/30/30	Optional
Optional	3km (fibre)	30/30/30	Optional
No	100/100/100	64/132/260	Optional
No	100/180/300	132/250/250	Optional
No	100/180/300	132/250/250	Optional
No	100/200/300	260/260/260	Yes
No	100/not known/not known	250/250/250	Yes

RMON = Remote Monitoring MIB RIRO = Ring-in/Ring-out PLL = Phase Locked Loop

Systems and SMC have eliminated the need to buy separate bridges by building bridge-routing intelligence into every network module on their latest hubs.

Greater fault tolerance is another hub advantage. Almost all modular hubs now on the market include features like redundant power supplies and hot-swappable modules (see Table 2 on page 92). If something does go wrong with the network, faults can usually be traced quickly because the hub serves as the central connection for the entire network. (With distributed MAU networks, troubleshooters often have to go from one wiring closet to the next to track down a problem.) Also, most Token Ring hubs accommodate modules that handle different types of LAN cabling (unshielded twisted pair, shielded twisted pair, coax and fibre).

Most modular hubs can handle different types of LANs, allowing net managers to put all their network connections in one box and manage them from one console. Lots of vendors sell Ethernet and FDDI slot-in modules. A few, including Cabletron and SynOptics, now offer ATM connections as well.

For added flexibility, some hub vendors now offer a feature called port switching (see 'Evaluating Port Switching Hubs' in the September 1993 edition). Port switching lets net managers assign ports individually to different segments on the backplane. Conventional modular hubs handle segment assignment on a per-module basis, which means all ports on a given module must be assigned to the same group.

Vendors now offering port switching include Bytex, Digital Equipment Corp, Optical Data Systems, SynOptics, and Xyplex. Initially much more expensive than conventional modular hubs, port-switching hubs are coming down in price.

High-end routers from vendors like Cisco Systems and Wellfleet Communications also can serve at the core of a collapsed backbone network. However, there are a couple of disadvantages to using them, one of which is that they tend to be very expensive.

Despite the advantages that modular hubs offer, they aren't an automatic choice for all Token Ring installations. In fact, modular hubs may be overkill in some networks. In particular, sites running low-speed applications — such as 3270 terminal emulation — that aren't yet experiencing performance problems can remain with the tried, trusted, and familiar MAU-based approach.

To Switch Is Divine

The continued use of MAU-based architectures in smaller corporate outposts, such as remote branch offices, is a clear indication that the tried-and-true approach continues to hold appeal. Still, demand also is building for higher-performance products. And it is that demand that has led to the development of a new class of equipment known as Token Ring switches.

Token Ring switches embody what really amounts to a third-generation Token Ring design, allowing the full 16Mbps LAN bandwidth to be dedicated to individual network nodes (see Figure 2 on page 90). Switch vendors say the main application for their products will be in replacing collapsed backbone hubs or routers, improving performance by cementing the Token Ring into even smaller LAN segments. Alternatively, Token Ring switches can be used to connect modular hubs.

Figure 2: The Token Ring Switch

Token Ring switches are an inexpensive way to segment networks into smaller rings, letting designers dedicate full LAN bandwidth to individual devices like servers or interswitch connections. Full-duplex technology now in development effectively doubles ring speed.

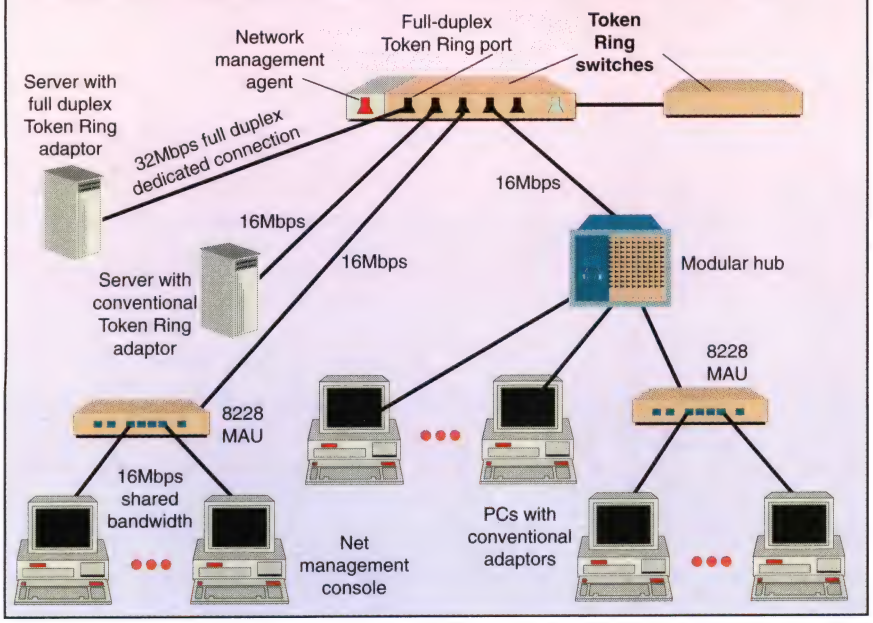

While Token Ring switches have the potential to allocate up to 16Mbps to every device on the network, vendors say that in the short term it is only servers that will have a pipe to themselves. PCs and workstations will share port connections via attachment to MAUs or hubs.

Regardless of what they are connected to, Token Ring switches are easy to install — net managers can use cabling and adaptors already in place. Three vendors now offer Token Ring switches: Cabletron, Israeli firm Netwiz, and SMC. When fully configured as Token Ring switches, these products are much more expensive than modular hubs but far less expensive than routers.

Vendors like IBM also are looking at ways to double the performance of switched connections, using a technique called full duplex. The technique, which works on dedicated switch connections to a single network node, involves disabling the token to allow the switch port and the adaptor in the attached device to transmit simultaneously, which enables data to be exchanged at 32Mbps. The IEEE recently started work on adding a full-duplex Token Ring specification to the 802.5 standard, and a stable draft is expected by the end of the year.

The coming year should see the arrival of another promising high-speed Token Ring alternative: products based on the nascent IEEE 802.12 100VG-AnyLAN standard. 100VG-AnyLAN defines a shared-media technology for LAN workgroups that can carry either Ethernet or Token Ring packets at 100Mbps. A stable draft of the standard should materialise this year. Hewlett-Packard plans to ship hub and adaptors implementing 100VG-AnyLAN this month, and while initial offerings will handle Ethernet traffic only, the vendor has pledged to add Token Ring support early next year.

100VG-AnyLAN initially had trouble picking up devotees during its development, but now support for the technology is building. Still, its future remains unclear because of competition with rival 100Mbps LAN standard 100Base-T. That standard, which is now making its way through the IEEE 802.3 committee, doesn't accommodate Token Ring.

Management Decisions

Choosing among the Token Ring building blocks is only one step in the design process. Planners also must decide which features available on today's Token Ring products are necessary for their network.

The single biggest product differentiator is management. Network managers must make a series of decisions when it comes to Token Ring management, each of which has an effect on cost and performance. The first question is whether to install manageable MAUs or hubs. If manageable devices are deployed, net managers must then choose a management protocol — either LAN Network Manager or SNMP. They also must decide where management intelligence should reside on the network — on every MAU and hub, or on selected devices only.

First-generation (MAU-based) Token Ring networks are managed using LAN Network manager software agents built into IBM-compatible bridges. These agents report diagnostic information about activity on remote Token Rings to a centrally located console. The main problem with this approach is that it gives the net manager only a limited amount of information about attached MAUs and hubs. For instance, bridge agents do not deliver per-port diagnostic information, and they can't perform security duties, such as disabling a MAU port to lock a network intruder out of the ring.

Deploying management intelligence at the hub or smart MAU gives administrators a greater reach into local segments. In this regard, IBM's LAN Network Manager protocol has one big advantage over SNMP: it supports a part of the 802.5 standard that allows it to interrogate the chip set on Token Ring adaptors to find out elementary information such as device addresses or to remove a given adaptor from the ring. This is a neat trick, because it requires no special software on the adaptor. In contrast, the only way most users of SNMP management can get the same information is to install adaptors that come with an SNMP agent.

Of course, SNMP's ace in the hole is that it's the de facto standard for LAN management — a fact that IBM has tacitly acknowledged with the release of the SNMP-based NetView/6000 product. IBM recently addressed the lack of SNMP support in its products with a new version of LAN Network Manager, called LAN Network Manager/6000. The latest version can convert management information from IBM's proprietary format into SNMP messages that can be passed to any SNMP platform.

Some vendors, like 3Com, have gone even further with SNMP support by implementing SNMP's remote monitoring management information base (RMON MIB). RMON allows administrators to draw together performance and status information about remote devices at a central management console. Note that some vendors offer only limited RMON support; full support requires the ability to handle all 10 of the RMON management sub-groups. Its only with a complete set that net managers can completely leverage RMON's ability to track network usage patterns.

To Pay or Not to Pay?

The price gap separating unmanaged MAUs and hubs from those that feature management intelligence can be wide. For example, in an unmanaged configuration, the Crossbow Plus from ADC Fibermux supports 168 Token Ring connections, at a cost of $195 per port (assuming a fully loaded box, with all ports belonging to the same segment). Making the hub SNMP-manageable requires the addition of a dedicated module, which costs $7,106. This decreases the maximum number of ports to 158 (the net management module comes with two ports, but takes up a slot in the hub that would otherwise be occupied by a 12-port

Continued on page 94

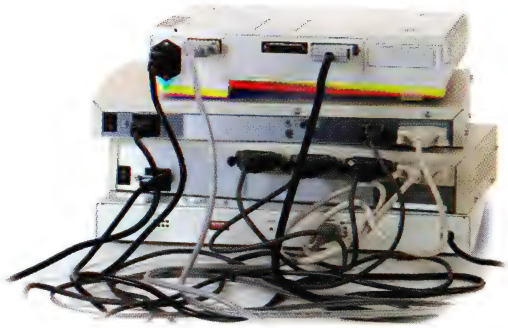

First Stackables.

Wellfleet INSIDE

Now Snappables.

Introducing SnapLAN™, the first remote office hub that anyone can install. And everyone can use.

Stackable hubs made remote LANs possible. Now Fibermux has taken the next step. We made them easy.

Just slide a SnapLAN module into the grooves and listen for the 'snap'. That's SnapLAN's revolutionary internal coupler making all the power and signal connections. Automatically.

No cables. No confusing backpanel wiring. No problems.

The use of flash memory in SnapLAN's integrated Wellfleet router and management module ensures remote operation stays just that. Remote.

Your office staff need never touch the SnapLAN box as all updates, diagnostics and maintenance are done from the central office.

And as SnapLAN can hold any combination of Ethernet or Token Ring module, your remote offices can remain as flexible as they choose.

SnapLAN. Solves remote connectivity and management problems, in a snap.

I need more information on SnapLAN's remote office solution. Please rush me literature on the SnapLAN snappable hub.

Title_____ Name_____ Position_____

Company_____ Address _____

_____ Phone _____ Fax _____

Please fax to: (02) 975 5337

ADC **Telecommunications Australia**
5B/4 Skyline Place, Frenchs Forest,
NSW 2086 Australia
Tel: (02) 975 1499

ISO9001 CERTIFIED

Table 2: A Guide to Modular Token Ring Hubs

VENDOR	PRODUCT	SLOTS	MODULE CONFIGURATIONS	BRIDGE OR BRIDGE-ROUTER MODULES
Ace/North Hills Dataplex (03) 210 3333	MAS 1000	12	8 DB-9 or RJ-45; 16 RJ-45	None
Andrew Corporation (03) 357 9111	Groupmaster/ Branchmaster	5 or 11	4 RJ-45; 8 RJ-45; 2 RJ-45 RIRO	Bridge-router planned
Bytex Corporation Techway (02) 975 2122	Intelligent Switching Hubs 7730/7760	6 or 17	12 RJ-45; 2 RJ-45 RIRO; 4 RJ-45 RIRO; 2 ST fibre RIRO; 4 ST fibre RIRO	Bridge-router
Cabletron Systems (02) 950 5900	MMAC-Plus	14	24 RJ-45	Bridge-router built into every module
	MMAC-M3FNB, M5FNB, M8FNB	3 to 8	12 RJ-45 or DB-9; 24 RJ-45; 10 RJ-45 with 2 DB-9 RIRO; 6 ST fibre; 12 ST fibre; 18 ST fibre	Bridge and router
Cameo Communications +1 (603) 888 8869	Ultrahub AH5000-12	12	12 RJ-45; 10 RJ-45 with 2 RJ-45 RIRO	None
Chipcom (02) 416 0653	Oncore Switching System (sold by IBM as the 8260)	17	8 RJ-45; 20 RJ-45; 2 RJ-45 with 2 FC, SMA or ST fibre RIRO	Bridge
	Online System Concentrator (sold by IBM as the 8250)	17	8 RJ-45; 20RJ-45; 2 RJ-45 with 2 FC, SMA or ST fibre RIRO	Bridge
Digital Equipment Corp. (02) 561 5252	DEChub 900 Multiswitch	8	8 RJ-45	None
ADC Fibermux (02) 972 1499	Crossbow Plus-FX6700	14	12 RJ-45; 10 RJ-45 with 2 ST fibre RIRO; 6 ST multimode fibre	None
Fibronics International Olivetti (02) 368 1988	GigaHUB XH150	12	12 RJ-45 with 2RJ-45 RIRO; 16 RJ-45; 2 RJ-45 RIRO; 2 fibre RIRO	Planned
Lannet Toren (03) 242 5050	Multinet LET-10, LET-20, LET-36	20 or 36	4 IBM-type, DB-9, or RJ-45; 10 RJ-45; 4 ST fibre; 2 DB-9 RIRO; 8 RJ-45 RIRO; 6 RJ-45 with 2 ST fibre RIRO	Bridge and bridge-router
Optical Data Systems Netland (02) 622 8505	ODS 1095/1096/1097	4, 7 or 12	32 RJ-45 or 50-pin with 2 ST or SMA fibre and 2 RJ-45 RIRO; 24 RJ-45 with 2 RJ-45 RIRO; 22 RJ-45 or 50-pin with 4 ST or SMA fibre RIRO; 18 RJ-45 with 2 RJ-45 RIRO; 20 ST or SMA fibre; 4 ST or SMA RIRO	Bridge-router
Plexcom ADE (03) 543 2677	CC8000 Series	1, 2, 4, 10 or 14	12 RJ-45 with 2 RJ-45 RIRO; 6 ST fibre	Bridge
Proteon (02) 955 8555	Series 90 Intelligent Hub	3 or 10	2 IBM-type with 1 IBM-type RIRO; 5 IBM-type; 10 DB-9 or RJ-45; 8 DB-9 or RJ-45 with 2 DB-9 or RJ-45 RIRO; 4 DB-9 with 2 ST fibre RIRO	Bridge-router
Racal-Datacom (02) 936 7023	Internext INX5000	3 or 13	10 RJ-45 with 2 RJ-45 RIRO; 16 RJ-45	Bridge
RAD Data Communications Dataplex (03) 210 3333	RADring	20	2 DB-9 or RJ-45; 2 SMA or ST fibre; 4 DB-9, RJ-45, RG-59 or RG-62	Bridge
SMC (02) 238 2206	Elite Switchinghub ES/1 Token Ring switch	5	1 DB-9 or RJ-45; 4 DB-9 or RJ-45	Bridge-router built into modules
SynOptics Communications (03) 853 0799	3000NT	12	12 RJ-45; 6 ST fibre; 12 RJ-45	Bridge or bridge-router
	5000NT	14	24 RJ-45; 20 RJ-45; 6 RJ-45 with 2 ST fibre or DB-9 RIRO	None
3Com (02) 959 3020	Linkbuilder Focus Series	2 or 5	2 RJ-45 RIRO; 8 RJ-45; 12 RJ-45; 20 RJ-45	None
	Linkbuilder MSH	11	12 RJ-45	Planned
UB Networks (03) 693 8200	Access/One	2, 5 or 11	12 50-pin with 2 50-pin RIRO; 22 50-pin with 2 50-pin RIRO	Bridge router
Xyplex NetArch (07) 393 1933	Network 9000	3, 6, or 15	12 RJ-45	None

*For 16Mbps networks, distances in metres except where noted 1 = Passive backplane RIRO = Ring-in/Ring-out NS = Not supported

AGGREGATE BUS SPEED	SNMP	LAN NETWORK MANAGER	MAX LOBE LENGTH CAT 3UTP/ CAT 5 UTP/STP*	MAX NODES CAT 3 UTP/CAT 5 UTP/STP	BEACON RECOVERY
16Mbps[1]	Via dedicated module	None	100/100/100	100/100/260	None
16Mbps[1]	Via dedicated module	Planned	100/100/185	132/132/256	Implemented in hardware
800Mbps/ 3Gbps	Via dedicated module	Via bridge-router module	75/150/150	72/200/250	Optional, implemented in software
10Gbps	Built into modules; RMON	None	200/250/300	150/150/250	Implemented in software and hardware
650Mbps	Via dedicated module	Via bridge module	200/250/300	150/150/250	Optional, implemented in software
72Mbps	Via dedicated module	None	50/200/200	132/132/132	None
13Gbps	Via dedicated module; RMON	Via bridge module	NS/100/175	0/132/250	Optional, implemented in software
2Gbps	Via dedicated module; RMON	Via bridge module	NS/100/175	0/132/250	Optional, implemented in software
3Gbps	Built into modules	None	44/65/100	256/256/256	None
264Mbps	Via dedicated module	Via dedicated module	60/100/150	256/256/256	Implemented in hardware and software
12Gbps	Basic built into module; advanced via dedicated module	Basic support built in; advanced via dedicated module	100/100/150	260/260/260	Optional, implemented in software
1.552Gbps	Via bridge or bridge-router module; RMON	Via bridge or bridge-router module	180/180/360	132/132/260	Optional, implemented in hardware
10Gbps	Via dedicated module; RMON	Built into modules	100/100/100	72/72/260	Optional, implemented in software
1.4Gbps	Via bridge-router module or dedicated module; RMON	Via bridge module	100/120/150	150/150/140	Implemented in hardware and software
52Mbps	Via bridge-router module; RMON	Via bridge-router module	45/175/250	100/260/260	Optional, implemented in software
1.3Gbps	Built into some modules	None	60/100/145	250/250/250	Implemented in hardware
16Mbps[1]	Via dedicated module	Via dedicated module	100/200/300	72/104/250	Optional, implemented in software
800Mbps	Built into chassis	None	100/200/300	72/72/255	None
382Mbps	Via dedicated module; RMON	Via bridge module	100/180/300	132/250/300	Optional, implemented in hardware
20Gbps	Via dedicated module; RMON	None	100/180/300	132/250/300	Optional, implemented in hardware
500Mbps	Via dedctd. module; RMON	None	100/200/300	260/260/260	Implemented in h/w and s/w
2.6Gbps	Via dedctd. module; RMON	None	100/200/300	260/260/260	Implemented in h/w and s/w
320Mbps	Via dedicated module	None	100/not known/not known	250/250/250	Optional implemented in software
300Mbps	Via dedicated module	None	100/150/100	128/128/128	None

From page 90

network module), increasing the per-port cost to $243.

The price difference is even wider if the Crossbow Plus houses multiple Token Ring segments — likely, given the trend to break rings into smaller parts to gain performance benefits. Each segment would require its own management module to get full statistical information about traffic on each ring. Some hub vendors significantly cut management expenditures by building SNMP intelligence into either their hub chassis or some or all of their network modules.

Net managers who decide to install manageable MAUs have two choices. First, they can buy from a vendor that includes management intelligence (again, LAN Network Manager or SNMP) in every MAU on the network. Second, they can opt for products from vendors selling complementary MAUs — one with management intelligence and one without. With these products, the management MAU collects data from the unmanaged MAUs and then forwards it to a central console. This approach, pioneered by IBM with its 8230 CAU/LAM products, has since been adopted by others — including ODS, Olicom and Proteon.

The basic idea behind the CAU/LAM (or master/slave) design is to spread the cost of management. Be warned though; most master/slave MAUs — including the ones from IBM and Olicom — restrict the reach of the network, since the management MAU and the unmanaged MAUs must be stacked in the same wiring closet (the MAUs are connected by short drop cables). The same restriction does not apply to products like the Elite Series 400C from SMC and the Linkbuilder TR Series Stackable Hub from 3Com. IBM says it is working on a LAM that can be remotely attached to a CAU.

Plexcom has a different take on the CAU/LAM architecture: it sells a dumb MAU, called the Plexring P8035SX, which can be remotely managed from a module installed in its CC8000 Series modular hub.

Another management feature to look for in smart MAUs and hubs is whether they can be managed out-of-band using a separate, low-speed connection. With these products, net managers can continue to get information about device status even when the Token Ring itself is beaconing — a condition that normally prevents management information from being passed. One hub vendor that offers both in-band and out-of-band management is Andrew Corporation.

The Jitter Bug

Along with management issues, network planners need to be aware of the ability of devices to handle problems that typically plague Token Ring installations. The two most common are signal attenuation (also known as jitter) and beaconing.

There are two basic causes of jitter. The first is near-end crosstalk and reflections caused by impedance mismatches in cabling. The second involves the Token Ring adaptor components that amplify and forward the Token Ring signal as it passes through each network node. Both problems lead to the same result — disruption to the timing of the Token Ring signal.

As jitter gets worse, nodes start to have problems reading incoming packets, leading to what are known as soft errors. These include read failures, very slow response times, and an inability to access a server. As the jitter increases further, some stations will start to drop from the ring, and when it gets really bad the network may crash.

Jitter is a cumulative problem — it gets worse in proportion to the length of the Token Ring and the number of nodes attached to it. Consequently, the conventional way to keep jitter in check is to limit the ring length and the number of nodes. And although jitter is present in all Token Ring networks, it is far worse on networks wired over UTP cable, which is much more susceptible to crosstalk than shielded twisted-pair cabling. That's the reason that the maximum number of nodes recommended for UTP Token Rings is 72, far less than the 260 recommended for LANs cabled over STP.

In the last couple of years, some vendors have started to include so-called active jitter attenuation components in their MAUs and hubs. These are intended to control jitter so that it does not affect network performance or limit ring size. Active products — as opposed to passive MAUs and hubs, which do nothing to control jitter — use one or more of three different methods to control jitter:

- *Active, non-retimed solutions*: In this, the simplest and least expensive approach, the MAU or hub includes an elementary electronic component that amplifies the incoming signal before forwarding it. No attempt is made to correct the timing of the signal.
- *Tank or ringer circuit retiming*: With this technique, a circuit containing capacitors and inductors amplifies the signal. Further, provided the signal is not badly distorted, the circuit will correct the signal timing. This provides a higher level of defence against jitter than active, non-retimed solutions, but it's also more expensive to implement.
- *Phase locked loop (PLL) relocking*: PLL also comprises components that amplify and retime the signal. Proponents claim it provides a higher level of protection against jitter than tank or ringer circuits. It is also the most expensive of the jitter control techniques.

Vendors include jitter reducing components in different places on their products. Some include them on all lobe and ring-in/ring-out ports on all of their slot-in network modules and MAUs. Others feature them on only some of the ports, and only on selected modules and MAUs.

Bigger Rings

The real-world advantage of using MAUs and hubs with active jitter-control components is that it enables designers to build longer rings with more nodes.

Some vendors say they can support hub-to-workstation distances greater than 100 metres. Bytex, for example, claims to accommodate 200 nodes at distances up to 150 metres over Category 5 UTP with its Intelligent Switching Hub 7760/7730. In practice, few users will want or need cable runs of that length. Research by AT&T Network Systems shows that 95% of users are located within 60 metres of a wiring closet.

3Com, which uses active components in its products, says another advantage is that it enables the mixing of different types of cable in the same Token Ring. Usually, impedance mismatches preclude net managers from using STP and UTP cabling in the same network.

Using active components also makes the task of installing MAUs a good deal easier. Net managers installing a Token Ring using IBM's 8228 MAU, or look-alike MAUs from other suppliers, must work out the maximum possible lobe length (the distance from the MAU or modular hub to the desktop) by taking into account a number of different factors, including the number of nodes on the network, the number of MAUs, and the distance between MAUs.

Despite the advantages of using active components, some question whether the extra cost incurred is money well spent. "A lot of the debate about active versus passive hubs is academic, since more powerful applications mean most people have decided to put fewer rather than more stations on their Token Rings," notes Kevin Tolly.

Another big problem with active components is that they are not covered in the 802.5 standard, raising the spectre of possible interoperability problems.

Beacon Beaters

Another crucial differentiator among Token Ring products is the way they handle beaconing, the most common type of problem on Token Ring networks. Beaconing indicates a break in the Token Ring. If a break occurs in a ring segment, the Token can't be passed around that segment, causing devices on the network to issue a beaconing signal. The three main causes of beaconing are improper attachment of adaptors to the ring, adaptor failures, and cable breaks.

Suppose a user plugs a faulty adaptor into a Token Ring. The problem adaptor cannot pass the network token, causing a ring break. The next adaptor along the ring waits for a predetermined period of time for the token to show up. When it doesn't, it starts to beacon — that is, it sends out a

stream of alarm messages indicating a problem on the ring. In theory, beaconing signals should travel all the way around the ring to the faulty adaptor, which then uses the 802.5 error-recovery algorithm built into the Token Ring protocol to recognise that it is the cause of the problem and automatically disconnect itself from the ring. In practice, if the adaptor is seriously flawed, it may be unable to read the beacon, and it will not disengage from the ring. As more stations fail to receive the Token, they start sending out alarms, ultimately crashing the network.

Vendors take several different approaches to deal with beaconing. Some go no further than relying on 802.5's self-healing mechanism. That's not as half-hearted as it might sound: figures from IBM indicate 10 out of every 11 beaconing conditions will not crash the ring if left to sort themselves out.

But clearly there are a couple of problems with this approach. One is that with the 802.5 error-recovery algorithm, it generally takes from 20 to 55 seconds to remove a problem station from the ring. But 20 seconds is an eternity on a Token Ring, and net managers can generally kiss their user sessions and SNA sessions goodbye.

Furthermore, if the error-recovery mechanism fails to work, network administrators typically have no choice but to send technicians to each MAU or modular hub to find the problem source. The diagnostic method usually deployed is to disconnect and reconnect each network node until the problem station is found.

A Sense of Speed

Some vendors are now building features into their products to alleviate the beaconing problem. A number of products now use rate detection to prevent users from plugging 4Mbps adaptors into 16Mbps Token Rings, or vice versa — an extremely common cause of beaconing. Vendors offer this speed detection facility by using components that sense the adaptor's data rate prior to its insertion in the ring. IBM has put a neat twist on rate detection: while its 8230 intelligent MAU cannot perform speed detection, IBM this year became the first vendor to ship an adaptor — the Auto LAN Streamer — that automatically detects the speed of the ring and adjusts its rate accordingly.

Some vendors, including Fibronics and 3Com, offer signal detection components that validate that the incoming signal comes from a Token Ring device. This prevents users from accidentally trying to plug other types of equipment into the network.

Rate and signal detection components prevent beaconing form occurring in the first place. Some vendors now offer features ensuring that any station causing a beaconing problem is removed from the ring faster than the 802.5 scheme can manage. Generally, beacon recovery is accomplished using algorithms incorporated as part of the network management facilities in their MAUs and modular hubs. To achieve rapid recovery times, some algorithms automatically expel the last station to hold the token, on the often-justified assumption that this is the one most likely to be causing the problem.

Beacon recovery and other state-of-the-art features can be a big help in keeping Token Rings up and running. But the availability of equipment with such high-end capabilities doesn't necessarily spell the end for Token Ring workhorses like the IBM 8228 MAU. In fact, net managers baffled by the variety of options open to them can take comfort in the fact that no one type of Token Ring equipment is ever likely to emerge as an outright technology winner. For every designer who opts for the performance and manageability of a modular hub, another will still prefer the simplicity of an 8228 (or look-alike) MAU.

Stephen Saunders is local area networks Editor for Data Communications magazine, based in California.

ACCTON

Just right for work groups

EtherHub 1000 System. Secure, Stackable and Intelligent

INTELLIGENCE

Accton's EtherHub 1000 Systems provide intelligent SNMP management solutions for various cabling applications.

Designed to Expand with your Business

Successful businesses grow, and to stay successful, you need a network that will grow with your business.

The Accton EtherHub 1000 System offers a choice of cabling, including Fibre Optic, Coaxial and twisted pair. Network administrators can easily expand the number of workstations on the network with support for up to 168 nodes. Add the *EtherHub-MGT* or *EtherHub-12mi* together with *AccView*™, our Windows-based SNMP network management software and achieve total control of each port.

For robust network security across several stacks add the *EtherHub-SBH*. This Secure Backbone Hub prevents unauthorised network access by providing security checks on port, segment, server and protocols. Its unique multi-segment design also provides a significant increase in network performance.

Cost Effective and Stackable

The Accton EtherHub 1000 System has been designed to save you money as your business grows. By starting small and adding components as you need them means your original network investment is protected.

The Accton EtherHub 1000 System offers flexibility at an affordable price. The 12-port *EtherHub-12i* is a perfect start for a small Ethernet network. Simply stack an additional EtherHub unit onto the system as workstations are added to the network.

ISO 9001 Symbol of Quality Excellence

Accton MPX cards include extensive driver support, station monitoring software and Novell's Universal Netware Client.

AC & E Sales
N.S.W.–Tel: (02) 712 3555 Fax: (02) 719 8548
VIC. –Tel: (03) 417 3622 Fax: (03) 416 0750

Anabelle Bits
N.S.W. –Tel: (02) 313 6155 Fax: (02) 313 6143
QLD. –Tel: (07) 862 2622 Fax: (07) 862 2616
A.C.T. –Tel: (06) 285 4616 Fax: (06) 285 2295

CTM Group
QLD. –Tel: (07) 257 1687
Fax: (07) 257 1689

Microbits
S.A. –Tel: (08) 362 9220
Fax: (08) 362 8445

Accton networking products offer the best price/performance in the industry.
Telephone or fax for a copy of the Accton Product Guide, the definitive guide to our complete range of Ethernet networking products.

Kevin Tolly and
David Newman

Putting the Squeeze on WAN Costs

Cross-country connectivity is expeditious but expensive. With compression ratios that tested as high as 7:1, remote Ethernet bridges put the big squeeze on wide area network costs.

Can vendors live up to their claims about data compression? When the Data Comm Test Lab checked out remote Token Ring bridges in June, we found a big gap between promise and performance (see 'Testing Remote Token Ring Bridges' in the June edition). Claims of 4:1 compression to the contrary, some of the bridges failed to hit 2:1 — even with highly compressible text. Could things be any better with Ethernet?

The answer, happily, is an emphatic 'yes.' Simply put, the six remote Ethernet bridges we put through their paces in our Test Lab can compress text by 3:1 or better. And some of this equipment actually deserves the much-overused label 'plug and play.' Considering that leased-line tariffs and management represent the two biggest costs of remote bridging, our findings are very good news indeed for network managers.

Two products merit special mention. The Series 3000 compression bridge/router from the Magnalink Division of Telco Systems compressed text by the astounding ratio of nearly 7:1 on the US-style 56Kbps lines we used for our tests. In fact, Magnalink's compression algorithm was so good that the Lab had to toughen up its testing for this round. The Remote Compression Bridge (RCB) from RAD Network Devices worked right out of the box — without us having to worry about configuration routines or other details. RND's bridge also delivered decent performance for a very competitive price. For those reasons, both the Magnalink and RND bridges earned Tester's Choice awards.

A third product, the LAN2LAN/MPR from Newport Systems Solutions, deserves an honorable mention. With its compression ratio of 6.7:1, it can put almost as much squeeze on data as the Magnalink bridge.

But Newport's cards and software approach means shelling out for a host PC, and that can put its price on the steep side.

The Test Lab also evaluated the MBR/6003 from Microcom, the 4810 LAN Bridge from Retix, and the NetBuilder from 3Com. Given that all the products we looked at are solid performers, net managers have plenty of options when it comes to remote Ethernet bridges.

What to Look For

Compression is very important for bridges using low-speed links — almost a given in today's enterprise — which is why we only evaluated products that can put the squeeze on data. But there are many other features that net managers must consider, such as the number of LAN and WAN ports, filtering options, maximum speed of the WAN link supported, and top data rate at which compression is possible (see Table 1 on page 98).

We also rated bridges on how easy they were to set up and how much work is really involved in making those common configuration changes. In addition, we judged how simple it is to set all-important filters, which are key to keeping unwanted traffic off costly wide area network links. And we took a look at the clarity of product documentation.

Table 1: Selected Vendors of Remote Ethernet Bridges

VENDOR	PRODUCT	MAXIMUM LAN/WAN PORTS[1]	WAN SPEEDS	MAXIMUM COMPRESSION SPEED	COMPRESSION TECHNOLOGY	FILTERS
Magnalink Division of Telco Systems Techway (02) 975 2122	Series 3000 Compression Bridge/Router	1/1	9.6Kbps to 10Mbps	2.048Mbps	Hardware and software	MAC, SAP, NetBIOS, SNAP, custom
Microcom (02) 410 9953	MBR/6003	1/2	9.6Kbps to 2.048Mbps	1 Mbps	Hardware	MAC, SAP, custom
Newport Systems Solutions LAN Systems (02) 901 3655	LAN2LAN/MPR	4/12[2]	1.2Kbps to 2.048Mbps	2.048Mbps	Hardware	MAC
RAD Network Devices Toren (03) 242 5050	Remote Compression Bridge (RCB)	1/1; second WAN port available for dial backup	4.8Kbps to 2.048Mbps	2.048Mbps	Hardware	MAC, SNAP, custom
Retix (02) 369 1333	4810 LAN Bridge	1/1	2.4Kbps to 2.048Mbps	128Kbps	Software	MAC, SAP
3Com (02) 959 3020	NetBuilder	2/2	1.2Kbps to 6.312Mbps	64Kbps	Software	MAC, SAP, NetBIOS,SNAP,custom

1. Maximum for model tested; more ports are available from some vendors.
2. 24 WAN ports are available without data compression.

MAC = Media access control
SAP = Service access point
SNAP = Subnetwork access protocol

Of all these issues, ease of use ranks highest. Remote bridges are typically deployed in branch offices. A network manager will usually make a field call to install the bridge, but after that any changes will either be made remotely (over the network) or be done by someone with little or no technical expertise.

For management from a central site, all the products tested come with SNMP agents that report to any SNMP network management console. This can be useful when monitoring status and gathering traffic statistics, but the ability to change bridge operating parameters via SNMP remains limited at best.

In fact, for all the talk about SNMP, remote bridges still rely on proprietary configuration packages. That's not a critical flaw, since all the bridges we tested feature serial ports for external modems and dial-up lines. Still, that's an extra step. Worse, we found some key features missing from many vendors' management utilities.

Five Difficult Pieces

To get a sense of how easy — or tough — it is to use these bridges, we identified five representative tasks: changing the maximum size of frames on the WAN; changing the WAN link speed; enabling or disabling data compression; setting a filter for one MAC (media access control) address; and altering a local MAC address. Then we worked our way through the list and rated all the bridges item by item on a scale of 1 to 10 (see Table 2 on page 100). We also took a look at each vendor's documentation, assessing it for clarity.

WAN frame size is critical on Token Rings, where one 4K frame — the legal limit — can choke a 56Kbps WAN pipe. It's not as important on Ethernet, although it takes only five 1,518-byte frames — the largest allowed — to slow a low-speed link enough to cause errors or session timeouts. That means that net managers *should* have a way to control WAN frame size, but of the bridges tested only those from Magnalink and 3Com meet that requirement.

Network managers also need a way to change the speed of the WAN link. Clocking is usually determined by external devices supplied by telcos or US-style CSU/DSUs (channel/data service units), but in some cases it's necessary to set it independently — when setting up backup WAN circuits, for instance. Newport and RND offer no way to accomplish this.

All of the vendors allowed compression to be turned on and off, something net managers may want to do when testing new devices or troubleshooting line problems. The best products in this regard were Microcom's and RND's whose software utilities allow changes to be made fairly easily. At the opposite extreme is Newport's bridge, which requires not just that the bridge be brought off-line but actually disassembled to disable compression.

On the LAN side, we determined how difficult it is to set a MAC filter and alter the LAN adaptor's MAC address. Newport was the clear winner in the first category: all that's necessary is to enter the MAC address of the device to be filtered. We should note, however, that remote Ethernet bridges generally support fewer types of filters than allowed by their remote Token Ring counterparts. For example, neither Newport nor RND offer SAP (Service Advertising Protocol) filters. That's a key omission for sites running Novell NetWare, because the huge number of SAP frames NetWare generates will eat up costly WAN bandwidth. And only 3Com's NetBuilder offers NetBIOS filters, which cut down on overhead on SNA internetworks and on those running Lotus Notes. Both types of filters, in contrast, are a regular part of the feature set of remote Token Ring bridges.

Only Retix and RND allowed LAN interface MAC addresses to be changed. This ability can come in handy when setting 'soft' addresses to conform with network naming conventions. MAC address changes on the bridge are a relatively minor issue, however, since end-stations rarely, if ever, address frames to the bridge.

Take The Money and Run

WAN tariffs and admin may account for the lion's share of remote bridging costs, but the bridges themselves still cost *something*. Once network managers have their products up and running, though, the thing that matters most is squeezing as much data as possible through the WAN pipe. This is where remote Ethernet bridges really shine: The top performers compress text data by ratios approaching 7:1, and even the slowest compresses text at better than 3:1.

Bear in mind that the compression achieved on real-world networks will probably be lower. These ratios represent best-case performance, using highly compress- ible text data, which is likely to be only part of the traffic mix (with the balance made up of less-compressible binary data). We also tested for worst-case performance, using highly random precompressed data, which cannot be squeezed down much further (see 'Test Methodology' on page 102).

The Ethernet vendors owe their superior performance to more advanced compression techniques, or even to a combination

SUPPORT FOR 10 INDIVIDUAL/10 GROUP FILTERS	NETWORK MANAGEMENT
Yes/yes	SNMP
Yes/no	SNMP
Yes/yes	SNMP; HP OpenView
Yes/yes	SNMP
No/no	SNMP
Yes/yes	SNMP

of techniques. A very good case in point is Magnalink's software-based compression, which operates on WAN links to 128Kbps. Its proprietary algorithm buffers incoming data and discovers patterns that can be represented by shorter strings. It also places data about frequently used strings in a dictionary, thus allowing it to learn frequently repeated patterns.

Magnalink also offers hardware-based compression on WAN links to E1 speeds (2.048Mbps). Hardware-based compression is usually faster, but Magnalink's compression software outperforms its hardware on 56Kbps links by about 70%.

Repeat Offender?

Magnalink's compression software, in fact, is so clever that it forced us to rethink our test methodology. We began this evaluation by offering each bridge a custom crafted string of 20 frames, repeated over and over. Since we used 512-byte frames, the total amount of test data was about 10Kbytes. That's not enough data to stress Magnalink's compression algorithm, as we discovered when our measurement gear reported compression ratios of 10:1 for precompressed data (since this data has already been compressed, the result should be about 1:1). Obviously, something was amiss.

The problem was that Magnalink's bridge buffered the entire 10Kbytes of data; then examined it and discovered that each 10Kbyte payload was identical to the one that preceded it; and instructed the remote bridge to keep repeating the data in its buffer. As a result, very little traffic actually traversed the WAN link.

Properly stressing Magnalink's bridge meant offering a payload big enough to overrun the buffer, so that it wouldn't simply be stored and repeated. Magnalink's engineers suggested that 80 frames, or about 40Kbytes of data, would do the trick. In reality, we needed to send 176 frames — more than 90Kbytes — before the bridge attained a throughput of roughly 1:1 when handling precompressed data.

Once we verified that we were offering enough traffic to obtain valid results, we proceeded to offer ASCII text with compression enabled. Even with the larger number of frames, Magnalink's product still delivered effective throughput of nearly 7:1 (see Figure 1).

In the interest of consistency, we retested the other bridges using the larger number of frames. Throughput results were unchanged except those for the LAN2LAN/MPR, indicating that Newport also uses some form of buffering to achieve its high compression ratios. Notably, the only vendors affected by the change in test traffic were the two top performers — Magnalink and Newport. The results presented here are the valid ones — that is, the results were derived from a payload large enough to stress each vendor's compression algorithm.

To establish a performance baseline, we also tested with compression disabled; all vendors were capable of operating at or near WAN line speed.

Latent Tendencies

Piling up frames in a buffer raises another concern — latency, or the amount of time each frame spends inside the bridge.

Gauging the effect of latency upon performance can always be tricky. A bridge with a very fast processor, for example, should boast an astronomical frame-per-second rate. But the very same bridge could well dump a large number of frames in a buffer before they're sent onto the network. the net effect of this sort of delay is a drop in effective throughput.

Compression and decompression, by definition, require bridges to do extra work — thus extending the amount of time that frames spend inside a bridge. it might be assumed the latency added by compression would serve to reduce throughput, but we found that isn't necessarily the case. In our Token Ring tests, in fact, there was a direct inverse relationship between throughput and latency.

Figure 1: Compression Compared

56Kbps WAN link, 512-byte frames

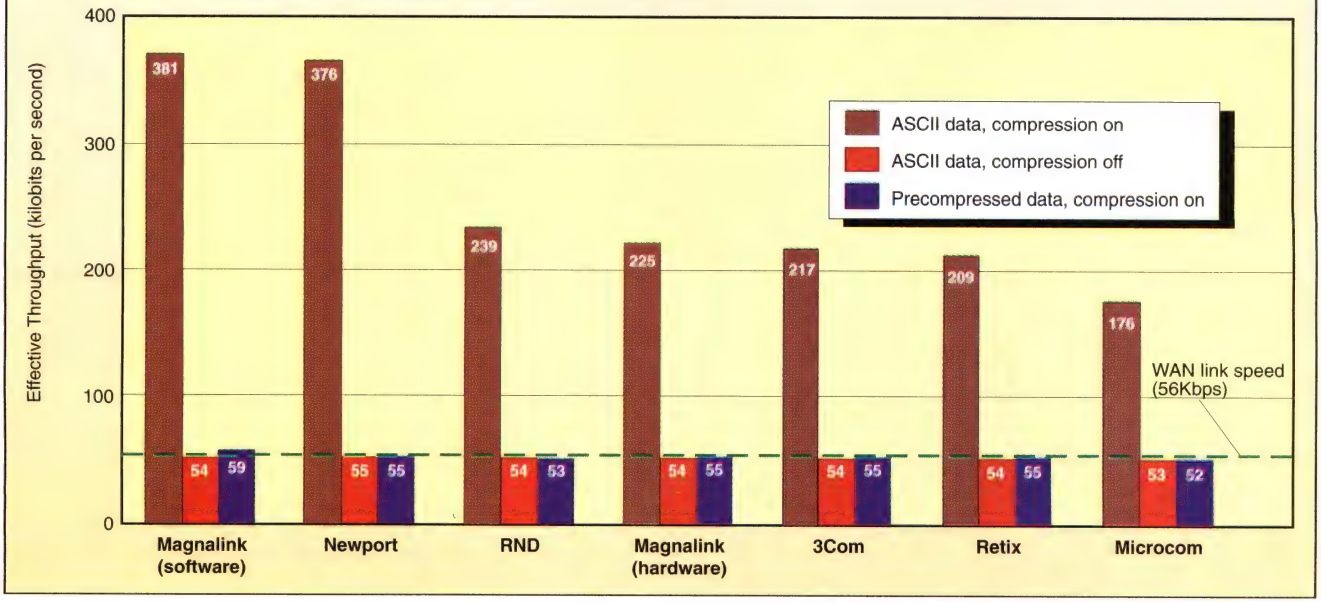

Table 2: Remote Ethernet Bridge Report Card

	MAGNALINK	MICROCOM	NEWPORT	RETIX	RND	3COM
FILTER CONFIGURATION						
Set filter to disable 1 MAC address	6	6	10	0	4	3
LAN CONFIGURATION						
Change MAC address of bridge's local interface	0	0	0	8	5	0
WAN CONFIGURATION						
Change maximum frame size on WAN side	6	0	0	0	0	7
Change link speed	7	7	0	7	0	6
Disable compression	7	8	2	7	8	3
SUPPORT						
Clarity of documentation	8	8	3	8	8	7

Ratings; Excellent, 9-10; Good, 6-8; Fair, 4-5; Poor, 1-3; Not supported, 0

This time around, there was not such a direct relationship; indeed, the vendors with the highest effective throughput — Magnalink and Newport — fell in the middle of the pack when it came to latency (see Figure 2 below). In both cases, large buffers added to these vendors' latency measurements.

The difference also can be explained by examining the type of data compression, hardware or software-based, that each vendor employs. Since computer chips process instructions faster than lines of code, hardware will always exhibit lower latency. But that speed doesn't always translate into a gain in effective throughput. Factors like buffering and variations in different vendors' compression algorithms make it possible to trade off higher latency for better throughput — a trade these vendors are willing to make.

Magnalink

The Magnalink Series 3000 Compression Bridge/Router not only proved to be the top performer in this test but also raised the standard against which the other products were measured. This strong showing isn't surprising, given that Magnalink specialises in WAN compression technology. The big shocker came when the bridge's software-based compression hit 6.8:1 with text on a 56Kbps WAN link, way out ahead of the 4:1 posted by the compression hardware under the same conditions. So much for the theory that hardware-based compression is always more effective.

The Series 3000 also did fairly well in our features evaluation. The management utility's command-line interface isn't as straightforward as the menu-driven offerings from Microcom, Newport, and Retix. But the documentation is clear and well-organised. And once the unit is up and running, it's easy to keep tabs on it — either via front-panel LEDs or from any SNMP management console. Magnalink supports virtually all major types of bridging filters, and the bridge also allows group filters to be set.

The Series 3000 we tested has one LAN and one WAN port, along with a serial port for dial backup. The higher-end Series 4000 line comes with multiple WAN ports.

Figure 2: Plotting Latency vs Throughput

56Kbps WAN link, 512-byte frames

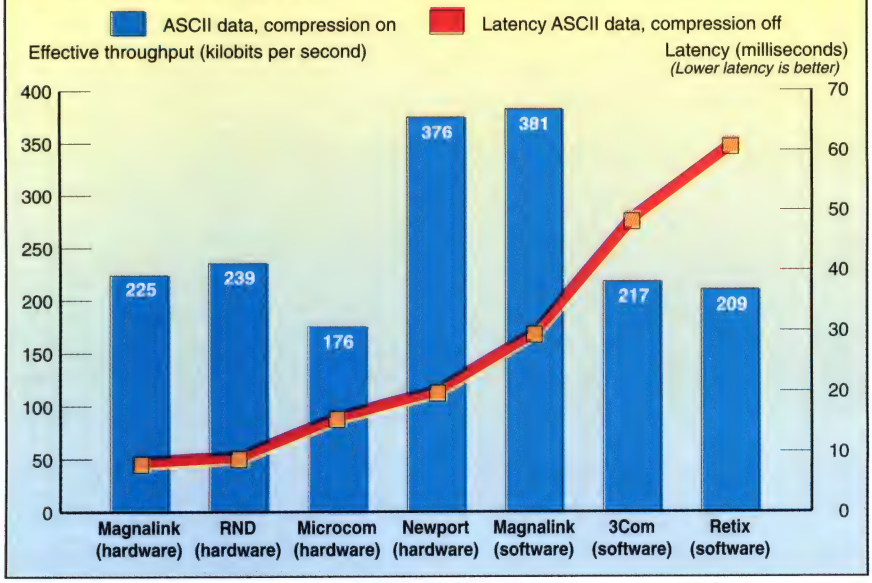

Microcom

We had a chance to check out the MBR/6003 in our remote Token Ring bridge test, when we found it easy to use and a very good performer at a good price. Microcom posted nearly identical results in this test; unfortunately, most of the other vendors did better.

In both the Token Ring and the Ethernet tests, the MBR/6003 compressed ASCII text by better than 3:1. While that was good enough to rank high among Token Ring

GUESS WHO'S NETWORKING IN PERTH?

Cabletron Systems, the world-leading manufacturer and marketer of local area network (LAN) and enterprise network systems, is committed to supporting it's Australian customers meet their networking requirements.

That's why we're expanding our Australian operation to add a new sales/support office in Perth to our existing offices in Brisbane, Canberra, Melbourne and Sydney (Head Office).

Australian customers can now access local technical or sales support through seventy employees dedicated solely to Cabletron service and support in Australia. And we have more than eight million dollars of inventory at our Australian headquarters to ensure you of fast, reliable replacement parts.

Cabletron prides itself on helping it's customers gain the competitive edge by meeting the demand for reliable and manageable information networks. Who do you call when you need the *Complete Networking Solution™?*

CABLETRON. THAT'S WHO.

CABLETRON SYSTEMS P/L

The Complete Networking Solution™

SYDNEY – HEAD OFFICE
Unit 8
Allambie Grove Business Estate
25 Frenchs Forest Road
FRENCHS FOREST NSW 2086
TEL: +612 950 5900
FAX: +612 950 5950

CANBERRA
Unit 2, Ferranti Building
Ground Floor
99 Northbourne Avenue
TURNER ACT 2601
TEL: +616 257 2122
FAX: +616 257 2726

MELBOURNE
11th Floor
580 St Kilda Road
MELBOURNE VIC 3004
TEL: +613 526 3639
FAX: +613 526 3673

BRISBANE
One Park Road
MILTON QLD 4064
TEL: +617 367 1750
FAX: +617 368 2923

PERTH
1 Havelock Street
WEST PERTH WA 6005
TEL: +619 429 8844
FAX: +619 429 8845

For further information on Cabletron's networking products, call your Cabletron sales representative or (02) 950 5820.

Test Methodology

Eleven leading vendors of remote Ethernet bridges were invited to participate in this test: Andrew Corporation, D-Link, Crosscomm, Hewlett-Packard, Magnalink Division of Telco Systems, Micom Communications, Microcom, Newport Systems Solutions, Retix, RAD Network Devices, and 3Com. Five vendors declined the invitation: Andrew, Crosscomm, D-Link, HP, and Micom.

Our test bed reflected the most common configuration for remote Ethernet bridges: two 10Mbps Ethernet LANs were connected over a wide area network link by a pair of bridges (see the figure).

We used a MOD-SW35 modem eliminator from US-based VIR to simulate the WAN link. A Novell NetWare 3.11 server was attached to one LAN, while a NetWare client was attached to the other LAN. (The NetWare client was used only to have the bridge learn its MAC address; it was removed from the LAN during actual testing).

To measure bridge throughput, traffic was generated on one LAN and sent across the WAN link to the other LAN where data rates were measured with a protocol analyser. Three kinds of traffic were used: first, we established a performance baseline by sending ASCII data with bridge compression disabled. Second, we sent the same ASCII data with compression turned on and noted the difference in effective throughput. Finally, we sent precompressed data through the bridge to verify that the product's compression algorithms would not degrade throughput by attempting to compress incompressible data.

To create the data, the Novell NetWare client copied a 1Mbyte file to the server; this file exchange was captured with an Expert Sniffer network analyser from Network General and then edited for use on a Framethrower frame generator from LANquest Labs. We used a stream of 176 unique frames for each of our tests, each 512 bytes long. Compressed data was created using PKZIP, a data compression program from Pkware.

Throughput was measured with an Expert Sniffer by dividing the number of bytes received by the time for the file transfer. In addition, a high-end DA-30 protocol analyser from Wandel & Goltermann monitored network traffic, allowing us to note the point at which frames were dropped.

We used The Tolly Group's tried-and-true single-frame test to gauge bridge latency: the Framethrower sends a frame across the remote link, a DA-30 timestamps the frame on the sending and receiving LANs. The difference between the timestamps is considered the latency.

Kevin Tolly and David Newman

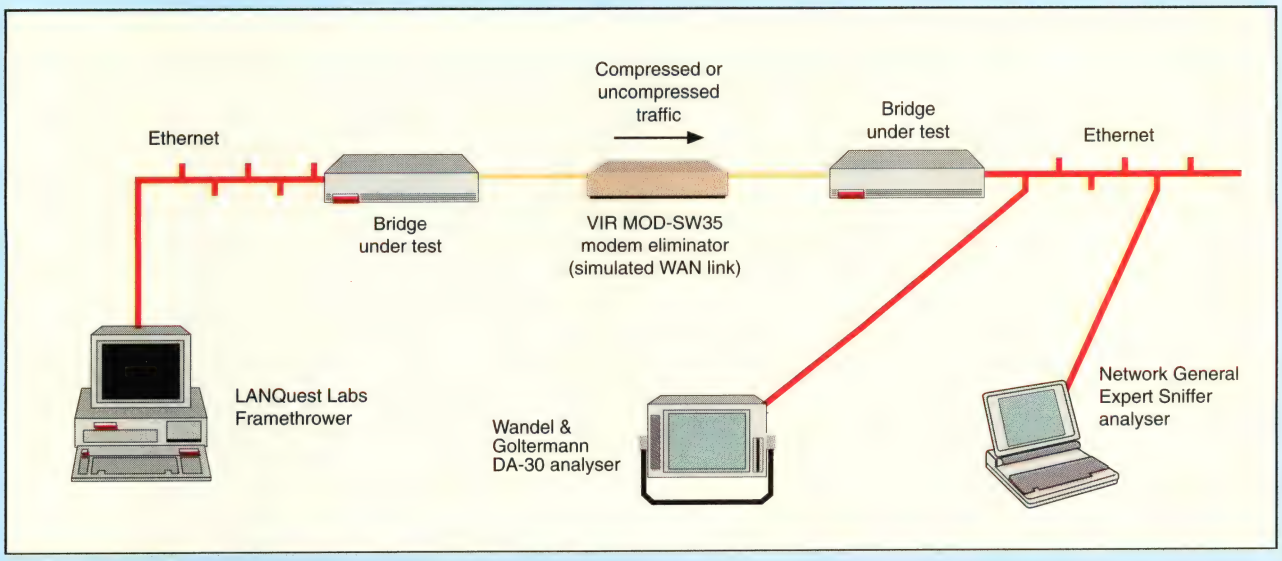

products, it means finishing dead last in the Ethernet tests. As mentioned, 3:1 compression of text is still very good, but with Ethernet bridges better performance can be had elsewhere.

The Microcom MBR/6003 remains relatively simple to work with. Although it's one of three bridge/routers we evaluated, its menu-driven configuration utility didn't bog us down with lots of routing options. Most routines were straightforward. The notable exception involves changing or disabling MAC addresses, which meant wading through several layers of menus. And as we noted in the Token Ring test, saving any change automatically causes the bridge to reboot without warning. That might be an advantage for non-technical users who don't know to reboot to effect a change, but for the experienced network manager the sudden reboot may come as an unpleasant surprise.

The product we tested was a modified PC with one LAN and two WAN ports.

Newport Systems Solutions

The LAN2LAN/MPR bridge/router delivers outstanding data compression and is easy to use.

Newport's package consists of a two-port WAN adaptor card, data compression daughterboard, and software — all of which can be installed in any PC with an ISA, EISA, or Micro Channel bus.

Configuration is through a series of menus, which we found extremely easy to navigate. We particularly liked Newport's approach to setting MAC address filters; all that's required is entering the MAC address in one pull-down menu. On the downside, not every feature displayed in the menus has been implemented yet. For example, disabling data compression appears as a menu choice, but to actually do so we had to remove the daughterboard from the WAN card. Newport says it is working to implement all features of its software.

The LAN2LAN/MPR compressed text at a ratio of 6.7:1, just slightly behind the 6.8:1 achieved by first-place finisher Magnalink. Like Magnalink, Newport supports data compression all the way up to E1. For network managers interested in gathering compression data, the LAN2LAN/MPR's management software even displays the bridge's compression ratio in real time.

Newport supplied its bridge/routers on a pair of Bravo LC 4/66d PCs from AST Research; the 66MHz 486 processors in

I N T R O D U C I N G

interWATCH™

Multitasking LAN Analyzers

Featuring:

Performance, Value, Flexibility, Ease of Use and Low Cost

- interWATCH™ gives you seven layer decodes with support for over 140 different protocols while you monitor. Solve the problem as it happens.
- interWATCH™ provides real time source routing decodes. See traffic congestion as it happens.
- interWATCH™ has the proven real time performance and elegance found only with true multitasking operating system.
- interWATCH™ gives you comprehensive statistics and filters to quickly isolate problems between partners.
- interWATCH™ provides a flexible, easy-to-use, fully customizable human interface to simplify troubleshooting.
- interWATCH™ has Ethernet and Token Ring PCMCIA solutions for notebook PC's, and ISA and microchannel solutions for other configurations.

When you compare the features, capabilities and simplicity, of interWATCH™ and then add our 20 to 50 % price advantage, it's easy to see that by any measure no other analyzer on the market comes close.

GN Elmi Australia Ltd, 9 Prospect Street, Box Hill, Victoria 3128. Phone: (03) 890 6677 Fax: (03) 899 5553.
Branch Office: Level 14, 33 Berry Street, North Sydney New South Wales 2060. Phone: (02) 959 1092 Fax: (02) 959 1992.
New Zealand Office: Leatham Electronics, Wellington. Phone: (04) 385 9409 Fax: (04) 385 7263.

GN Elmi

these PCs helped boost Newport's standings in the test, but the zippy hardware adds to the price of each bridge. In addition, at least one Ethernet LAN adaptor is required.

RAD Network Devices

Sure, RND's Remote Compression Bridge (RCB) delivers better than 4:1 compression of text. Sure, it's at the low-end of the price scale. But those impressive feats pale besides RND's real achievement: the RCB works as soon as it's plugged in, making it a true plug-and-play device.

As might be expected, the RCB asks very little of users in the way of configuration. It automatically senses WAN link speed and assumes compression should be enabled. The RCB automatically cuts over to a dial backup port if there's trouble on the leased-line port (and it allows net managers to set a threshold for cutover).

The RCB furnishes three kinds of filters: those based on MAC address, on user-definable conditions, or on SNAP (subnetwork area protocol) traffic. Conspicuous in its absence is support for SAP filtering, which is a key way of keeping Novell Netware broadcast traffic off the wide area. Other features include SNMP support and data compression at any speed to E1.

The RCB is the bargain box of this round of tests. While it lacks some features found on more expensive bridges, it represents excellent value for the money.

Retix

Like RND, Retix offers a low-cost remote bridge that's easy to manage. But the 4810 LAN Bridge lacks a couple of features found on products from RND and others, such as support for high-speed data compression and the ability to set filters that cover groups of users.

Configuration of Retix's bridge is handled through a menu-based setup utility that we found easy to navigate. In fact, only the RND bridge was easier to set up. The 4810 also can be configured through Retix's proprietary central-site network management software.

The Retix bridge is equipped with one LAN and one WAN port; there's no dial backup option. The bridge's software-based compression can only be used with WAN links to 128Kbps.

Our performance tests indicate that the 4810 is capable of compressing text by a ratio of nearly 4:1 — a figure that would have been impressive in our earlier round of Token Ring tests but qualifies Retix only for a sixth-place finish here. Retix says it is working hard on optimised compression algorithms.

3Com

The NetBuilder offers more features than any other product in this test — full routing functionality, support for all major filter types, and complete integration with the rest of 3Com's line of bridge/routers. But those advantages must be weighed against the box's middling performance.

As we noted in the Token Ring tests, 3Com's menu-based configuration routines are of necessity more complex than those of simple products: after all, there are more options available. But we lost our way more than once and wonder whether all those functions should be shielded from users in remote offices. Fortunately, any NetBuilder product can be configured remotely.

The NetBuilder compressed text by nearly 4:1 on 56Kbps links. But the vendor's software-based compression can only work on lines operating at 64Kbps or less.

Kevin Tolly is President of the Tolly Group, based in Manasquan, New Jersey. David Newman is Testing Editor for Data Communications magazine.

Why not get the WHOLE story on communications?

Whether it's GSM, SDH, FDDI, OSI, ATM, ISDN, SNA, Frame Relay, Fastpac, Token Ring or Ethernet, you need to be on top of the rapid changes now taking place in the communications industry.

Each month in the pages of Australian Communications the region's top communications journalists, industry figures, academics, professionals and experts bring you the latest news and analysis on the technical, political and management issues facing this dynamic industry.

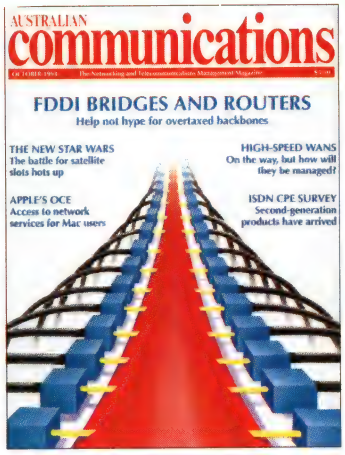

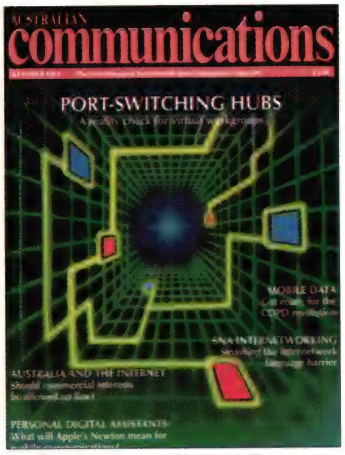

RULES AND CONDITIONS OF ENTRY

1. Instructions on how to enter form part of these Rules and Conditions of Entry. Entry is free and open to all residents of Australia (valid only in States where legal) other than employees of the promoters: Auscom Publishing Pty Ltd ACN 003 606 102, 4/541 Kent Street Sydney NSW 2000; SingCom (Australia) Pty Ltd ACN 002 864 897; their families and associated agencies.
2. To enter use the coupon provided. Only one entry per subscription is permitted. Each correct entry will go into the prize draw and the winning entry will be the first drawn on 1st December, 1994, by a representative of SingCom (Australia) Pty Ltd at 323 Castlereagh Street, Sydney, NSW. Entries not completed in accordance with the rules and conditions, or received after the closing date of 30th November, 1994, will not be considered.
3. The judge's decision is final and no correspondence will be entered into. The prize is not transferable nor redeemable for cash. All entries remain the property of *Australian Communications* magazine. The prize winner will be notified by telephone and written letter and their name published in the February 1995 issue of *Australian Communications*.
4. The winner will receive a trip for two consisting of return airfare to Singapore flying Singapore Airlines (from the winner's home state capital), and 5 nights twin share accommodation at the Westin Stamford. The prize is valid to December 1995. Accommodation and flights are subject to availability. The prize does not include meals or spending money. Prize value $2,918.00.

SUBSCRIBE NOW

☐ 2 years (22 issues) for $84 (save 30%) ☐ 1 year (11 issues) for $42 (save 30%)

Name .....

Title .....

Company .....

Address .....

..... Postcode .....

Telephone .....

I enclose a cheque for $ ..... ☐ Bill my company

Authorising signature .....

Charge my: ☐ Amex ☐ Bankcard ☐ Diners ☐ Visa ☐ Mastercard

Account No ..... Expiry date .....

Signature .....

NSW Permit No. TC 94/3359. VIC Permit No. 94/2316 issued 22/7/94.
ACT Permit No. TP 94/1387. Northern Territory Permit No. NT94/1103.

sing

POST TO: Australian Communications, 4th Floor, 541 Kent Street, Sydney NSW 2000
FAX TO: (02) 264 2244

ISDN Networking

3Com has announced the release of a comprehensive range of ISDN products for connecting remote and personal offices.

The new offerings include routers/bridges, adaptors and remote access servers. Products will be available for all major international ISDN standards, and they will operate under 3Com's Transcend network management system.

The company will provide ISDN Primary Rate Interface (PRI) and Basic Rate Interface (BRI) systems for the remote office (100 users or less) and personal office (less than 10 users).

For personal office connectivity, 3Com will offer a slide-in ISDN module for its AccessBuilder remote access server. The module will come with four ISDN BRI ports, and ISDN ISA and PCMCIA will also be available. For very small offices, the company will offer a Personal ISDN bridge, and to support larger numbers of personal offices 3Com will scale its ISDN offering using an AccessBuilder with an integrated PRI port, offering ISDN support for up to 30 concurrent users with over 256Kbps throughput each.

For the remote office market the company will offer its NetBuilder II router with dual ISDN PRI ports and its NetBuilder Remote Office bridge/router with one ISDN BRI port.

3Com (02) 959 3020

Branch Office Router

CTEC has announced Gandalf's 5250i Intelligent Branch Office Router, which has been specially developed for use in IP and IPX environments.

CTEC officials said the product features 8:1 data compression, providing up to 76Kbps throughput over 9,600bps lines and up to 512Kbps throughput over a single ISDN 64Kbps B Channel. Based on a 960CA RISC processor, the product features intelligent RIP spoofing, only transmitting changes in RIP across the WAN and saving on bandwidth usage.

Officials said that the 5250i's ability to automatically initiate connections on specific IPX or IP addresses or Time of Day provides savings on branch office connectivity by minimising line requirements.

The unit also supports over 2,000 routing entries and multiple subnet support, and can be used to sub-divide networks, create firewalls and provide alternate routes and redundant paths. It supports SNMP MIB II, comes with 1MB of Flash RAM for easy upgrade, and is priced at $5,195.

CTEC (02) 975 4722

The Axis NPS 630 can achieve throughput of up to 2,000Kbps

TR Print Server

Swedish firm Axis Communications has announced the release of a pocket-sized Token Ring print server which allows any PC on the network running NetWare, LAN Manager, LAN Server or TCP/IP, or any combination of these, to access any printer on the network.

The Axis NPS 630 and NPS 632 incorporate Axis's own 32-bit RISC processor, and are able to achieve throughput of up to 2,000Kbps. Both versions have a set of eight fully programmable logical printers, allowing automatic ASCII-to-Postscript conversion, string substitution and character set conversion, officials said.

The units also feature the ability to print from a multi-host environment, the selection of multiple bins/forms sizes, and the running of different printer emulations.

The Axis NPS 630 and NPS 632 will be available in September, priced at $795.

Intelligent Technologies (02) 891 6010

EISA ATM Card

Danish manufacturer Olicom has introduced the first of a planned family of ATM products, the Olicom ATM EISA Network Interface Card (NIC).

The new card is a 155Mbps SONET multimode fibre card for EISA PCs and servers which can be used to provide direct desktop connection for applications needing high bandwidth, such as imaging. Available immediately in beta form, final versions of the product will be available at the end of the year.

Olicom's ATM range, which was developed in partnership with Fujitsu, will include 51 and 155Mbps cards, and will sup-

Olicom's 155Mbps SONET multimode fibre ATM card for EISA PCs

Cray's approach avoids many problems of mixed SNA/LAN networks

New Cray Tunnel SNA Processor

Cray has made another addition to its host communications product range, the Model CS6100 SNA Transport Processor. The new offering integrates 'incompatible' LAN and SNA networks, protecting companies' substantial investment in SNA equipment and applications.

The Transport Processor is an integrated solution that provides all users with access to all network services. It sets up multiple tunnels (up to 100 per unit) one level higher than the Transport Level. Cray officials said this solution overcomes problems associated with conventional tunnelling on the Network Level, such as SNA protocol limitations, session time-out, local termination, SNA prioritisation and SDLC-to-LLC conversion. The Transport Processor handles SNA protocols through bridges and routers, providing SNA users with the same transparent access to the internetwork as LAN users.

Officials added that the product's multi-protocol approach offers several benefits: a single backbone can carry both SNA and non-SNA traffic, cutting costs; adding more protocols is simplified; and new technologies can be deployed easily by upgrading or replacing routers.

The SNA Transport Processor comes with a local management program for diagnostics and tracing in both SNA and OSI environments. A Cray Network Management Centre configures, manages and monitors, and stores and distributes software for each unit, and SNMP management is available through the SNMP agent running on the interface board.

Cray Communications (02) 451 6655

Cayman's branch office products feature flexible, modular design

Cayman Branch Office Solution

Cayman Systems has announced wide area connectivity support for its *i Series* family of workgroup routers. The GatorRoute iR and GatorStar iHR feature modular design and offer two flexible serial WAN ports for connection between remote office and corporate sites via leased line, analogue dial-up or digital dial-up services such as ISDN or frame relay. The products also incorporate two Ethernet segments for local routing at the branch office, and the GatorStar iHR includes an integrated 24-port 10Base-T hub with individually managed and monitored ports.

Protocols supported over the WAN link include IPX, TCP/IP, AppleTalk and DECnet, and leased line and dial-up links are supported via expansion cards. Two WAN serial cards are initially available, the Serial V.35 and the Serial RS232. The cards support both asynchronous and synchronous modes, and operate at speeds up to 115.2Kbps. The V.35 card can be used for either dial-up or leased line connections, while the RS232 card is designed for dial-up or dedicated routing for modem or leased line connections. Both connect to CSU/DSU for leased-line WAN access or to a third party conversion device for ISDN and frame relay.

The PPP-compatible *i Series* can be managed with Telnet or Cayman's own GatorKeeper, and out-of-band management is available through the serial console port.

ADE Network Technology (03) 543 2677

port standard bus architectures such as ISA, EISA, PCI and PCMCIA.

Force Technology
(02) 971 1000

The Memorex Telex 5038 prints at up to 16 pages per minute

Network Printer

Memorex Telex has released the new 5038 Laser Printer, a 16-page per minute multiconnectivity printer for the 3270, AS/400 and PC printer markets.

The printer accepts integrated LAN attachment cards for both Token Ring and Ethernet environments, and features multiple connectivity options with coax, twinax, serial and parallel attachments and dynamic host sharing capability.

The Model 5038 has a 500-sheet input bin and optional additional 500-sheet bin, automatic emulation switching to facilitate automatic switching between ports, an optional duplexing feature and envelope feeder, and full compatibility with AFP (Advanced Function Printing)/IPDS (Integrated Printer Datastream), and IBM's SAA printer standard.

The Memorex Telex Model 5038 is available for $8,500 including tax, and comes with a 12-month on-site warranty.

Memorex Telex (02) 805 5805

New HP Omnibook

Hewlett-Packard has added a new model to its Omnibook range of notebook PCs, the Omnibook 530. Officials said the new model offers a threefold increase in performance over previous machines, and also includes VGA-out for display flexibility and an optional floppy-disk drive.

Weighing 1.36kg, the system combines the fastest performance with the lightest weight of any 486-based subnotebook, according to officials.

In its standard configuration, the system features a 33MHz Intel 486SX CPU, 4MB of RAM (expandable to 12MB), a 130-MB hard disk and an additional PCMCIA Type II slot, and is priced at $3,140 including tax.

The Omnibook comes preloaded with MS-DOS 6.2 and Windows 3.11, HP's appointment and phone book software, and Travelling Software's Laplink Remote.

In addition, the system offers standard Omnibook features such as 'instant on,' which restores the files and applications being used when the machine was turned off; a full-size 85-key keyboard; a built-in pop-up mouse; and long battery life which provides up to four hours of continuous use.

Hewlett-Packard 13 13 47

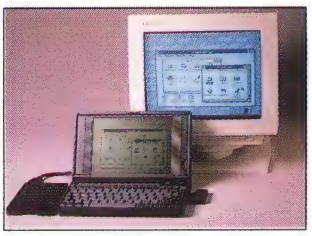

HP's new Omnibook is the fastest and lightest model on the market

Intelligent ISDN Range

Racal Australia has launched a new range of intelligent ISDN products to meet the demand for cost-effective routing and high network service levels.

The range includes the new DAP4000 Series of ISDN Basic Rate Terminal Adaptors, the DAP4500 Digital Leased Circuit ISDN Backup unit and the ISX5010 Access Multiplexer.

The DAP4100 provides access on demand to basic rate ISDN services for data terminal equipment (DTE). The unit is supplied with a range of interfaces including V.24, V.35, V.36 and V.21, and can be configured for speeds from 2.4 to 64Kbps, as well as in synchronous or asynchronous formats. Up to two DTE ports are available, which allows users to activate two simultaneous connections to two different destinations.

The DAP4500 provides resilience to a Digital Leased Circuit (DLC) in case of leased line failure, continuously monitoring the connection and providing automatic back-up over

Racal's new ISDN product range can handle high network demand

Perle's 494E connects users to AS/400s at up to 128Kbps

Basic Rate ISDN in the event of a DLC failure.

The ISX5010 Access Multiplexer incorporates a flexible architecture enabling users to add voice and data channels and higher speed trunks and a full network management system.

The unit also has the capacity to connect a wide variety of terminal equipment remotely over a single or multi digital trunk, and this can be backed up by an integral ISDN terminal adaptor, maximising the efficiency of link and providing a high level of resilience.

Racal Australia (02) 936 7023

New AS/400 Remote Controller

Canadian vendor Perle Systems has released a new remote controller. The 494E remote controller allows net managers to connect users directly to one or more IBM AS/400s via SDLC network, V.35 or X.21 host interfaces, with speeds of up to 128Kbps.

Fully compatible with IBM's 5494 controller, the 494E offers a range of enhanced features like extra Twinaxial and Token Ring user support, multi-session mode and multi-host capability.

The unit supports a maximum of 28 Twinaxial devices and 80 Token Ring users for 108 device connections. In Enhanced Perle mode, the unit can support a maximum of 160 Token ring and 112 Twinaxial devices simultaneously using the six available card slots. Officials said this is an increase of 164 device attachments over the same product from IBM.

Enhanced mode also provides 'virtual' multi-session and multi-host support for Twinax and Token Ring users. Twinax users can increase the number of sessions up to 1000 virtual sessions simply by adding memory to the Perle 494E.

The Token Ring card for the Perle 494E is fully compatible with all IBM Token Ring products including MAC-layer bridges. When fully configured for 160 users, a Twinaxial card is not required, providing users with a cost-effective Token Ring-to-AS/400 gateway.

Pricing for the 494E starts at $7,990 (including tax) for the 28-device Twinaxial model.

Perle Systems (02) 416 0650

Router Management

Retix has announced its new RetixVision graphical SNMP-based router management tool, which supports all popular Unix network management packages including HP OpenView, IBM NetView/6000, SunNet Manager, and NetLabs/DiMONS.

The new software simplifies device configuration and speeds the pinpointing and diagnosis of

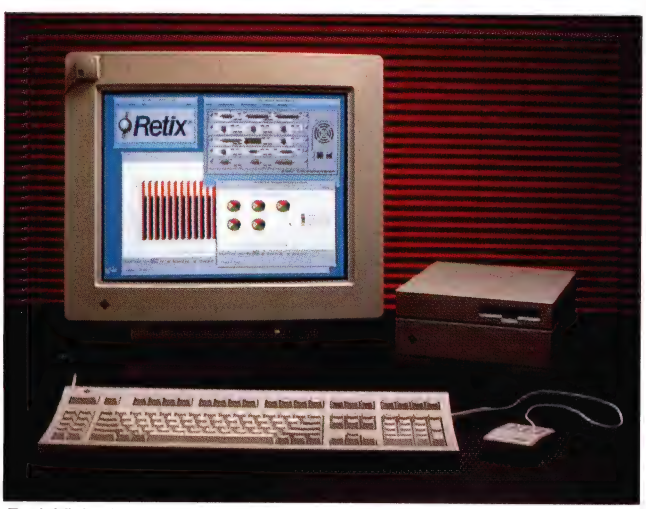

RetixVision's intuitive interface makes Retix router configuration easy

problems in enterprise networks, officials said.

RetixVision uses a intuitive user interface to allow network managers to easily and accurately configure all Retix devices, located either centrally or remotely. It features graphic depictions of each device's connector panel, an on-line help facility that provides access to all MIB object definitions, graphic display of device statistics and error conditions, and easy retrieval of performance, status and configuration data.

RetixVision can run in standalone mode on any Unix workstation, can be used in conjunction with other Unix-based network management systems, or can be combined with Retix Network Management Packs. It supports all Retix internetworking products including RouterXchange 7000 central site routers and 4000 Series boundary routers and bridges.

Retix (02) 369 1333

Lantronix's Ethernet transceivers boast performance at a low price

Miniature Transceivers

Lantronix has announced two new low-cost miniature Ethernet transceivers.

The model LTX-TA can connect twisted-pair cables in 10-Base-T networks, while the LTX-2A is designed for thin coax 10Base-2 networks. Officials said despite the units' low price tags they offer all most-needed features, and come with a five-year warranty.

The LTX-TA has automatic polarity correction, LED Power and Good Link indicators, and an RJ-45 network interface. The LTX-2A has a single LED Power indicator and a BNC barrel connector media interface. Both feature a switch-selectable heartbeat function.

The units measure just 6cm by 4cm by 2cm, come in rugged moulded cases and can withstand operating environments of up to 70°C.

PowerCorp (02) 476 3466

PCB BERT

Marconi Instruments has released a new handheld PCB Bit Error Rate Tester (BERT) for in-service and out-of-service testing of 32-channel systems.

The new 2840 2Mbps Transmission Analyser is a portable network tester which provides a range of performance test and diagnostic capabilities normally reserved for larger, more expensive systems, officials said.

The 2840 weighs only1.3kg, has a large backlit LCD graphics display and a user-friendly interface, and allows up to eight sets of test results to be stored, as well as up to 10 full instrument set-up configurations.

Marconi's Transmission Analyser is a lightweight, easy to use unit

A 'print to RAM' capacity of 10,000 events is provided for storage and test results for analysis and print-out at locations remote from the test site. The unit's battery life covers a full working day, and batteries can be quickly recharged from AC supply or a vehicle battery in two hours.

The Model 2840 also features a combined pattern generator and error detector, providing 2.048Mbps framed and unframed testing. Capabilities include in-service and out-of- service testing of E1 systems with

contiguous and non-contiguous n x 64Kbps testing and G.821 error performance analysis.

A highly accurate propagation delay measurement of 1-bit resolution enables loop-backs to be accurately and simply located. Digital signal frequency and level measurement facilities simplify analysis of signal degradation, while output level attenuation provides level margin testing information.

A full set of alarm lamps is provided on the instrument panel to provide user feedback on the status of the current measurement, and the unit also provides error distribution and alarm histograms to aid interpretation of results. Results can be analysed down to one-second resolution using the display's zoom capability, streamlining the pinpointing of network interference problems.

The 2840 comes with an autoconfigure mode, and can also be configured remotely via an RS-232 remote control interface. The autoconfigure mode detects the presence of framing, signalling and CRC, and channels are scanned to detect 64-Kbps test patterns.

Marconi Instruments
(02) 638 0800

Local Router

Gandalf International has released the LANline 5250L multi-protocol router, which the company says is a secure and robust product designed to link internal LANs while overcoming incompatible protocol problems.

Using an Intel 75Mips RISC processor, the LANline 5250L offers 12,000 packets per second throughput when routing, and up to 14,880pps when bridging, officials said.

The unit also simplifies migration from old to new Ethernet frame types by providing virtual network support via automatic Ethernet frame translation and support for four Ethernet frame types (Ethernet 2, 802.3 Raw, 802.2 LLC and 802.2 SNAP).

The unit has a Windows-based console for easy configuration set-up, and provides a range of statistics for use in maintenance and diagnostics.

Gandalf Pacific
(02) 437 5977

Network Management Utilities

Symantec has released an integrated system of network management tools and utilities that allow end-user applications and data to be managed efficiently and cost-effectively.

The Symantec Network Series comprises: Norton Administrator for Networks, Norton Anti-Virus for NetWare, Norton Desktop for Windows 3.0, Norton Enterprise Backup, Network Menuing Admin Pack 2.2, Norton DiskLock Administrator, Norton pcAnywhere and Norton Utilities Administrator.

Officials said the suite of products can dramatically increase network manager productivity by centralising administration tasks and automating routine, time-consuming jobs.

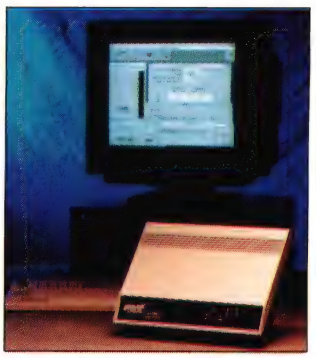

The LANline 5250L deftly mixes a range of network protocols

One of the newest additions to the suite is Norton Utilities Administrator, the first automatic PC and disk management tool to operate over a network from a central console. The system's benefits include: automatic operation; centralised control of file system recovery and the ability to repair each workstation from a central console; the ability to create custom Rescue

disks; seamless integration with the Norton Administrator for Networks; and the inclusion of The Norton Utilities 8.0.
Symantec (02) 879 6577

Frame/Signalling Analyser

Wandel & Goltermann has introduced the PA-41 Frame/ Signalling Analyser, a new field service instrument designed for commissioning, maintaining and troubleshooting 2Mbps circuits. A combined frame and channel associated signalling analyser, the PA-41 has two receivers (2Rx), allowing both directions of a 2Mbps circuit to be monitored simultaneously.

Wandel & Goltermann's PA-41 troubleshoots 2Mbps circuits

Circuit testing can be performed with framed or unframed signals, and using the unit's built-in CAS simulator, complete call sequences can be programmed for testing switching through analogue exchanges.

Other features include drop and insert functions of testing of multiplexers/demultiplexers, digital cross connects and automatic protection switches. The PA-41 also includes an interface for memory cards, and measurement data can be stored on a RAM card and processed on another PA-41 or on a PC.
Wandel & Goltermann
(03) 690 6700

Portable Test Set

Wedgetail Technologies has announced the availability of the new Schlumberger Themis family of telecommunication data analysers, which the company says are the first products that can be configured for both SDH and PDH network testing, or for any combination of the two.

The Themis can be configured for both SDH and PDH network testing

The analysers feature a small footprint and an open VMEbus-based architecture to provide modularity and flexibility, with a consistent user interface across all test applications.

Officials said the new range facilitates the testing of hybrid network environments by offering PDH testing at all hierarchy levels from 64Kbps to 140-Mbps, synchronised with SDH analysis at STM-N levels.

In addition to operating in a fully transparent mode, the transmission and reception parts of Themis can also be programmed independently, allowing PDH or SDH equipment and networks to be tested separately.

The system's modularity lets users start testing on any hierar-

Acsys's compact MicroConnect bridge suits small workgroups

chical level and progressively add functionality to meet new testing requirements as network topologies change. Starting at the basic Themis mainframe, SDH options can be added to provide 155.52Mbps STM-1 with or without 622.04Mbps STM-4 test capability, and with optical and electrical interfaces. PDH options can also be added for 64Kbps to 140Mbps test capabilities, including analogue and all intermediate digital hierarchical levels.

From a single connection point, Themis can measure different levels of the signal simultaneously: the SDH B1/B2/B3 transport parity bytes, the mapped PDH signal in the SDH container, and the different PDH tributaries in the PDH signal. Standard access, protected and high impedence connections simplify interfacing to the equipment or network under test.

Themis is able to qualify all significant transmission parameters conforming to the latest ITU-T M2100 and G826 Recommendations, from bit, code and frame to Bi/CRC/FEBE blocks. Integrated signal frequency, propagation delay time and network synchronisation measurements can eliminate the need for additional test equipment, officials said.

Bit and block error rate testing, detailed timing analysis and extensive disturbance, error and alarm generation ensure that transmission equipment and networks can be fully qualified and monitored, including tributary multiplexing, PDH/SDH mappings, protection systems, alarm thresholds and synchronisation limits.

The unit features a user-friendly interface and autoconfiguration, and a macro programming feature enables complex test sequences to be set up and stored.

Wedgetail Technologies
(02) 415 3944

Small Footprint Bridge

Acsys has announced its new MicroConnect local bridge, a compact, ultra-high-speed Ethernet bridge suited to small workgroup applications.

Company officials said the plug-and-play MicroConnect is the smallest bridge with full Ethernet functionality on the market, measuring just 10.8cm by 5 cm by 2.2cm.

It connects directly to a hub's AUI port, eliminating the need for a separate AUI cable. It is unit-powered from the AUI device, and therefore does not require an external power supply.

Used for segmenting workgroups from the network backbone, the bridge can connect workgroups of up to 256 users, and provides full-wire speed filtering and forwarding of data at 14,880 packets per second, with a throughput delay of less than 15 microseconds. It also features automatic address learning and an ageing function to store repetitive traffic routes for quick routing.

The MicroConnect has two LEDs to provide feedback on bridge and network status. It features a 100-year mean-time-between-failure, and is sold with a lifetime warranty.

Acsys (02) 299 2799

Comms Software

Australian firm Stallion Technologies has recently released X.25 communications software, known as Xpac-25. The package provides a range of X.25

services for the interconnection of Unix-based systems over X.25 Packet Switched Data Networks.

Facilities on X.25 networks can be accessed by interfacing applications to Xpac-25's Application Programmer's Interface (API) or from character terminal sessions to remote hosts using an X.3/X.28/x.29-compliant PAD. Xpac-25 can also be combined with the company's Xtend-Router software to run applications supporting the TCP/IP socket library over an X.25 WAN.

Other facilities supported by Xpac-25 include Unix file transfer utility support (UUCP, ftp and rcp), and an X.25 specific fast file transfer feature which speeds file transfer and cuts costs by utilising the built-in error-checking mechanisms of X.25.

Stallion has also released the Xconnect dual-channel adaptor, which can be used for a range of communications tasks. The adaptor provides a synchronous communications channel using V.24 or X.21 signalling, with a second optional channel, and when used with Xpac-25, enables both links to operate at up to 19.2-Kbps.

Stallion also has available a Unix-based synchronous driver which allows users to write their own software drivers for Xconnect to handle specific communications tasks, such as the routing of TCP/IP packets over ISDN links.

Stallion Technologies
(07) 870 4999

Full Duplex Ethernet Switches

ADE has announced the new 10Base Switch Full Duplex Ethernet Switch, which can substantially increase data throughput, according to officials.

The product increases network capacity and server availability by providing a seamless, well-partitioned environment between Ethernet segments.

Using the new switch, clients and servers can conduct two-way multiple 'data conversations' at speeds of up to 20Mbps, whilst also preserving full compatibility with the existing installed Ethernet base.

Based on technology which allows each switch port to serve up to 1,024 users, the product works with full duplex cards from major vendors including IBM, Compaq and Digital.

ADE Network Technology
(03) 543 2677

LAN/WAN Optimiser

Techway has announced the new Series 5000 LAN/WAN Optimiser from Magnalink Telco Systems. The product is an advanced compression device which can increase the throughput of a bridge, router, front end processor or channel extender an average of two to four times, depending on the type of data.

The advanced hardware pipeline and multiprocessor architecture of the system minimises compression/decompression latency and maximises throughput, reducing bandwidth requirements and reducing congestion on the WAN, Techway officials said.

The system provides one or two input ports, to be connected to the WAN port of a bridge, router, gateway or SNA controller, and one main and one backup output port are provided for connection to a modem or multiplexer. Interfaces supported include V.35, X.21 EIA-530 and RS232C.

The Optimiser supports WAN speeds from 9.6Kbps to 2.048-Mbps (depending on the model) and buffers are built-in to ensure maximum throughput with no packet loss.

The system guarantees reliable compression and decompression of data, even in the presence of line errors, by implementing an LAPB fully automatic error detection and correction feature.

Techway (02) 975 2122

Which publication is read by more communications professionals?

ATUG - The Exhibition and Conference of the Australian Telecommunications Users Group is unquestionably the major communications event in Australia. Visitors to the annual event are widely considered to be the region's top communications professionals, responsible for the lion's share of its networking and telecommunications purchasing.

Upon registering at ATUG'93 visitors were requested to indicate the communications/networking publications they regularly read. Judging from the **4,690*** completed questionnaires *Australian Communications* is clearly the choice of communications professionals

" I regularly read...."

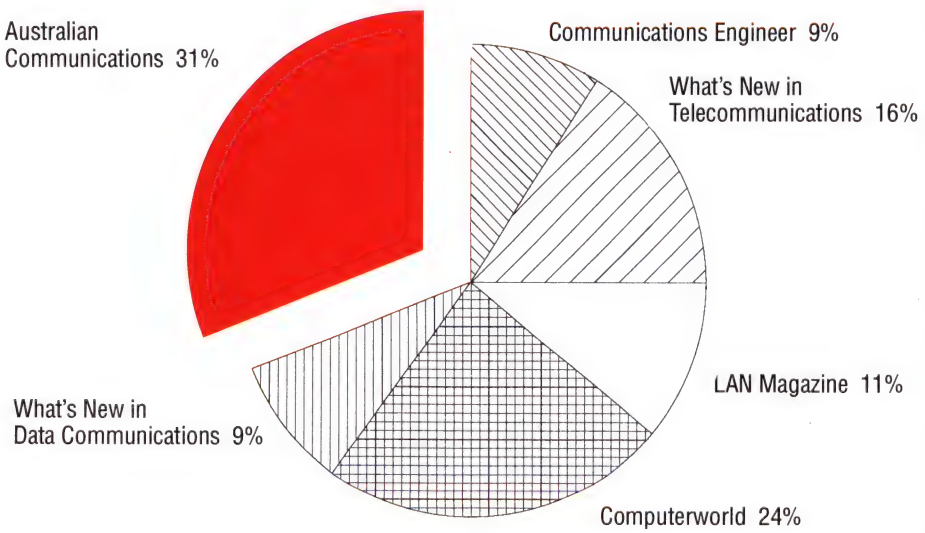

Data supplied by Riddell Exhibition Promotions, organisers of ATUG'93

* Does not include conference delegates, pre-registrations, or visitors with completed coupons from advertisements in *The Australian* and *Australian Communications.*

AUSTRALIAN communications
The Data and Telecommunications Management Magazine

International Advertising Representatives

Head Office
AUSTRALIA & NEW ZEALAND
Australian Communications
4/541 Kent Street
Sydney NSW 2000
Australia
Ph (02) 264 2200 Fx (02) 264 2244

USA & CANADA
M&T International Marketing
Miller Freeman Inc
411 Borel Drive Avenue
San Mateo CA 94402
Ph (415) 358 9500 Fx (415) 358 9739

EUROPE
A-Z International Sales
70 Chalk Farm Road
London NW1 8AN
United Kingdom
Ph 71 284 3171 Fx 71 284 3174

JAPAN
Media Sales Japan Inc
Tamuracho Building
3-3-14 Shimbashi
Minato-ku Tokyo
Japan
Ph 3 3504 1925 Fx 3 3595 1709

TAIWAN
Acer TWP Corp
4F 19-1 Lane 231
Fu Hsing N. Road
Taipei 104
Taiwan
Ph 2 713 6959 Fx 2 715 1950

KOREA
Jes Media International
5th Floor Banpo Building
40-7 Chamwon-Dong
Seocho-Gu Seoul 137-030
Korea
Ph 2 545 8001/2 Fx 2 549 8861

HONG KONG
Acer TWP Corp HK
Unit 2, 5th Floor
82-84 To Kwa Wan Road
Kowloon
Hong Kong
Ph 764 3857 Fx 764 3857

SINGAPORE
Eastern Publishing Associates
1123 Serangoon Road #03-01
Singapore 1232
Ph 296 6166 Fx 298 7551

ISRAEL
Nora Lifton
Kfar Saba 44444
Israel
Ph 52 970065 Fx 52 953194

CONSULTANTS CORNER

ATUG AUSTRALIAN TELECOMMUNICATIONS USERS GROUP

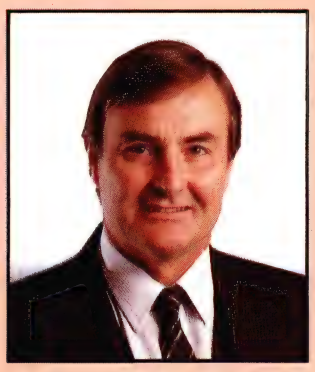

Editorial by Wally Rothwell

Pricing Rules Set to Change

Over the last year we have seen long distance call prices go down by about 18% and much improved quality of service and concentration on customers' needs, which we can attribute to the benefits of competition in the market.

But before we become too euphoric, we should realise that virtually all other tariffs have been at a standstill for the last 12 to 18 months. Part of the reason is that competition has not occurred in many areas — Telecom has proposed a one cent rise for local calls, payphone charges have gone up to 40 cents, and Optus hasn't made any big reduction in what many see as one of the most expensive satellite operations in the world.

So where there is still virtual monopoly in provision of many services, there appears no inclination by the carrier duopoly to bring tariffs closer to cost, and I'm not about to blame them for this marketing reality.

But another reason for tariff stagnation has been the uncertainty about the validity of some of Telecom's bundled tariffs, such as its Strategic Partnership Agreements and Flexi-plans.

ATUG

SYDNEY
Level 11, 80 Alfred Street
Milsons Point, NSW, 2061
Australia
Tel: (02) 957 1333, (008) 22 6281
Fax: (02) 925 0880
MELBOURNE
Level 1, 84 Johnston St
Fitzroy VIC 3065
Australia
Tel: (03) 416 1848
Fax: (03) 417 3002

In my June editorial, I referred to the recent amendment to the *Telecommunications Act* which clarified the operation of the dominant carrier pricing rules and gave Austel the necessary power to disallow anti-competitive tariffs.

In particular, the amendment meant that a carrier in a position to dominate a market should not be permitted to bundle services where that would be anti-competitive. An example of that activity would be the bundling of uncompetitive basic access and local calls with competitive services such as long distance. So Austel's task is to examine all bundled tariffs filed with the regulator and to decide if they are anti-competitive.

In order to do this, it was directed to develop a Decision Making Framework (DMF) and it is in the process of doing so, in extensive consultation with the major players, including ATUG. It has now had three public meetings, the third held after this editorial went to press.

In general ATUG has been supportive of the direction in which Austel is moving. Our position is that we need to be able to see some logic in the tariffs we are offered, and we want to unveil the ambiguity behind some of those tariffs. For instance, it's a bit hard to understand why, for the same pair of copper wires, there is a different access and service charge for business and private telephones, telegraph lines, Datel, DDS, DMS, Centel, ISDN, private voice lines, Austpac and Faxstream. One can understand a varying 'service' charge, but why should there be a differing access charge for what is really the same service?

This is one of the issues with which Austel seems to be coming to grips by suggesting that it will assess the three basic parameters of: the customer connection, switching functions and inter-exchange transport to assist it in its DMF. If the carriers are obliged to put a price on each of these functions, we the users will be able much better to understand the make-up of the tariffs we are expected to pay.

For instance, this may be of great benefit to the less advantaged, who have difficulty in paying for a telephone service. If they simply want service within their own local call area, they could be asked only to pay an access fee comprising the customer connection plus the switching function. Similarly, all those services with varying access and service charges would have a single access charge, and the service charge would differentiate the service. This process of clarity ought to lead to lower, more cost-based tariffs eventually.

This is the difficulty for Austel; it must find a way through the muddy swamp that has been the basis for our telecommunications tariffs until now and come out of it with a firm foundation for clearly calculated, competitive and comprehensible tariffs for the future. Of course, the marketplace will continue to define what we are prepared to pay, but the new Austel system will need to allow us to understand better what we are paying for.

Austel's aim is to have its DMF in place by the end of this month and, barring a legal challenge from one or even both of the carriers, that goal seems achievable.

Depending upon the various services offered by service providers and upon the regulations affecting the carriers, each has an important stake in the outcome of the DMF and, of course, each is lobbying hard with Austel for its own best outcome.

For the user, the outcome has to be that the tariffing regime must be fair, understandable, predictable and pro-competitive. It should also cause tariffs to move closer to cost. We think Austel has these same aims at heart and, for that reason, we have continued to support the thrust of its work in the DMF so far.

Wally Rothwell
Executive Director

ATUG's Extra Service

Members may be able to reduce their telecommunications costs by up to 30% following the introduction of a new savings scheme by ATUG, the 'ATUG Extra Corporate Discount Plan.' Under this scheme, ATUG is able to provide discounts on telecommunications services and other corporate services.

Telecommunications

⇒ Access to discounts via Telecom reseller, Network Exchange;

⇒ Aggregation of all members volume on Optus national and international AMPS mobile and BusinessNet services; and

⇒ Vodac GSM services.

Examples of the savings which users may be able to achieve via this scheme are:

⇒ For a one-person professional office, with combined phone, fax and mobile costs of about $200 each month, the savings under the ATUG Extra plan would be in excess of $600 per annum;

⇒ For a 13-person industry association, with combined phone, fax and mobile costs of $1,300 per month, the savings would be approximately $3,200 per annum;

⇒ For a small business with 20 personnel, and monthly telecommunications costs of around $2,800, the savings would be around $7,500 per annum;

⇒ For a 200-staff subsidiary of an overseas multinational, with a monthly bill of around $14,000, annual savings can be achieved in the order of $27,000.

The above discounts are, in the main, greater than the organisation can achieve by itself.

Travel

Westend Travel, the corporate incentive specialists, are providing ATUG members with a range of discounts, including:

⇒ Domestic airlines — 5% discount over and above all fares including 'special' fares (conditions may apply to special offers);

⇒ International airlines — 9% discount off recommended retail;

⇒ Better than corporate rates at more than 160 hotels and motels around Australia; and

⇒ Super deals on Budget car rentals.

Contact: David Campbell, Tel: (02) 891 5098.

Wine

For your chance to enjoy premium quality wine at a discount price, the Hunter Valley Wine Society has agreed to provide complimentary membership to ATUG members (normally $20).
Contact: Lois Etchell, Tel: 008 022 861

Equipment

Inchcape Office Products will give members a 25% discount off the recommended retail price of their full range of facsimile and copier products.
Contact: Lorraine Turtle, Tel: (02) 415 9444.

The savings under the ATUG Extra plan are in addition to existing ATUG membership benefits like discount rates to the ATUG convention and ATUG seminars.

ATUG Legal Line Service

Members should take advantage of ATUG's Legal Line Service. The service is run by ATUG director and chairman of ATUG's Legal sub-committee, Gerald Wakefield.
Members with legal queries regarding international telecommunications law and regulatory policy; radiocommunications; broadcasting and satellite law and policy; computer technology law; and intellectual property law are advised to seek this communications legal advice service by calling ATUG on (02) 957 1333.

Interception & Recording?

Want to find out more about call interception and recording? Or any other topic in the Telecommunications Broadcasting or Radiocommunications legislation? Then you need the EIS Electronic Law Book on disk. The new EIS Electronic Law Book has been released and now includes all the updates to the telecommunications legislation along with Broadcasting and Radiocommunications legislation. The publication, on a floppy disk, is available to ATUG members at the special price of $150 including — a saving of $50. A subscription including a further two updates is also available to members for $295 — a saving of $100.
To ensure that you have the legislation at your fingertips at all times contact the ATUG secretariat for details of how to order. Telephone (02) 957 1333.

A Towering Dilemma

As mobile carriers seek to honour the contractual obligations they incurred under the licence agreements purchased from the Government, they find themselves fighting a battle, not between themselves so much as on the public relations front. This battle is the result of anger and frustration felt by residents and councils alike and vented through the newspapers as carriers build mobile telecommunications towers with supposedly little if any consultation with the community.

The fine balance required between logistics, regulations, commitments to stakeholders and being a good corporate citizen is one that is testing carriers and is likely to continue to do so for a few years. It is worthwhile therefore, to take a look at some of the main factors which have to be taken into account as mobile networks develop.

The Towers

Dotting the horizon of most suburbs and becoming increasingly prevalent are 'steel trees' or structures required for mobile telecommunication networks. The term 'towers' to describe these structures conjures up a much taller structure than the terms pole or base stations which, by contrast, are used by the carriers. The majority of existing 'towers' are those installed for the establishment of Telecom's MobileNet AMPS network. The newer ones, and the ones causing greatest concern, are those required by Vodafone and Optus for the rollout of their respective GSM networks. Vodafone appears to be the carrier featured in many of the newspaper articles.

'Towers' or poles generally average around thirty metres in height with some reaching just on fifty metres. The output power of the transmitters is generally twenty five watts resulting in an effective antenna output of around two hundred watts. Depending on the terrain, a GSM base station provides coverage over an average three kilometre radius. Coverage can range from as low as one kilometre to as high as thirty kilometres. AMPS, on the other hand, averages around 10 kilometres coverage per transmitter. Network design and performance constraints, rather than propagation, limit the coverage. AMPS, like GSM, can be used to the limits of line of sight; some AMPS base stations cover an area of 120 kilometres or more. Whilst existing buildings and structures are used to mount the antennas where cost and physical constraints permit, free standing base stations are being used more often by all of the carriers. Free standing base stations often have space at the ground equivalent to a small house block to contain signalling and switching equipment.

Obligations

To meet the obligations under the licence, a carrier is required to achieve coverage to 80% of the population by the end of 1996. MobileNet's base station population of 920 supports their AMPS network for around 1.1 million customers. GSM, as a rule of thumb, is between 2.5 and 3.0 times more efficient when it comes to spectrum efficiency, however, the design of a GSM network differs considerably from AMPS and will in fact require more base stations to cover the same area. This means we will be looking at a lot of towers, poles, base stations or whatever they are called. This is especially true since geographical coverage is only one factor in the number of sites; others include network design, power output level, antenna configuration, filtering, interference and topography. It is estimated that each of the three carriers will be installing some 600 base stations.

Regulation

Prior to the deregulation of the telecommunications industry, Telecom was exempt from most State and Territory laws, including laws relating to the use of land and the protection of the environment. With deregulation came the *Telecommunications Act* and the Government's decision that carriers should, where practical, be subject to State and Territory laws.

There were, however, to be some exemptions such as the installation of telecommunications facilities, trenches and transmitters. Reports recently indicate however, that legal advice suggests that the exemptions do not extend to Commonwealth Law. As such, compliance with the Australian Capital Territory's National Capital Plan by the carriers is required.

These exemptions were to encourage development and were to be covered by a separate Telecommunications National Code, this Code to be applied by industry regulator, Austel. Until June 30th this year, the code had been in draft form, however, in February this year, the Department of Communications and the Arts sent to the Attorney General documentation for a final code. Negotiations had continued between the parties, to ensure that changes necessary to meet legal obligations did not conflict with policy requirements. The completed document was gazetted and tabled in Parliament on 30th June.

The Code now requires carriers, when establishing network facilities, to comply with certain guidelines which include:

- The *Environmental Protection Act, 1974*
- Standards Association of Australia standards for non-occupational exposure to radio frequency emissions
- A requirement to consult with relevant local councils and State government authorities regarding proposed facility installations.

According to the last of the above requirements, activities are not permitted to commence until:

- 28 days after notification to the relevant authorities; or
- 42 days after notification, if a relevant authority requests further information or suggests alterations to the proposal.

Under the Draft Code, regardless of the above notifications, councils and government authorities had no right of veto over the carrier's proposal. The code now forces the carriers to consult with the Department of the Environment, Sport and Territories when a dispute exists.

Co-location

Under the new Code, carriers are required to use their best endeavours to co-locate tower facilities. According to Telecom, the majority of their GSM antennas are co-located on AMPS towers. Whilst co-location seems a reasonable request, there are issues that must be considered, particularly for Optus and Vodafone. Some of these include:

- To prevent radio interference, transmitters need to be physically separated from each other. This may mean an increase in height by some five to ten metres. An increase in height alters the geographic coverage profile and may not give the desired quality of service in hilly areas such as across Sydney.
- The AMPS network will cease to exist by the year 2000. It is a possibility that Telecom will seek to reclaim the hardware and sell it overseas to a country at an earlier

stage of development. This would result in altered propagation characteristics and the probable need to relocate, or at least retune, the GSM transmitters and antennas on the site.

- For Telecom, co-location is fine because it has its network connection already in place. For Optus and Vodafone, co- location may mean an unwanted interconnect cost into the Telecom network. For Optus and Vodafone, a direct line of sight positioning gives them the chance to use microwave links back to their own network to minimise variable costs associated with a call.

For the Government, the problem is how best to balance the various elements. On the one hand residents (the voters) and councils (the persuaders) are protesting; fortunately for the politicians these protests are so far being directed at the carriers. On the other hand, rollout obligations cannot be reduced, or the carriers will choose to reduce service levels in the lower profit (regional) areas.

Dishes

Whilst most controversy to date has been focused on the towers, with Pay TV likely to invade our senses in the not too distant future we are likely to see a proliferation of satellite dishes of around a metre in diameter ultimately, covering most rooftops.

Those who have travelled abroad will be aware that the sight of towers, microwave links and satellite dishes is not pretty.

We all want Australia to have a telecommunications infrastructure. Mobile Communication is the key to such a capability; the trade-off, however, is the structures needed to provide the network.

ATUG

SYDNEY
Level 11, 80 Alfred Street
Milsons Point, NSW, 2061
Australia
Tel: (02) 957 1333, (008) 22 6281
Fax: (02) 925 0880
MELBOURNE
Level 1, 84 Johnston St
Fitzroy VIC 3065
Australia
Tel: (03) 416 1848
Fax: (03) 417 3002

ATUG'94 Conference Papers on Disk

All delegates who attended the ATUG'94 Conference will receive a copy of the conference papers on computer disk. Papers cover:

- ⇒ Telecommunications Update Seminar
- ⇒ Conference streams and workshops
- ⇒ ATUG Interoperability Network
- ⇒ The Open Learning Seminar

If, for some reason, you were unable to attend ATUG'94 and would like to purchase a copy of the papers on disk, please mail or fax the order form below and forward to:

By mail: Australian Telecommunications Users Group, P.O. Box 357, Milson's Point NSW 2061

Or fax: (02) 925 0880

Please forward _____ (qty) copies of the ATUG'94 Conference papers on disk at $200 for each copy ($300 non-members) for which I enclose payment of: $_____

Please tick payment method (DO NOT send cash):

☐ **Cheque** ☐ **Mastercard** ☐ **Bankcard** ☐ **Visa**

Cheques should be made payable to ATUG.

Credit card number:

☐☐☐☐ ☐☐☐☐ ☐☐☐☐ ☐☐☐☐

Expires: _____ Cardholder's Signature: _____

Cardholder's Name: _____

Company: _____

Address: _____

_____

_____

ATUG in South Australia

ATUG members in South Australia should take advantage of the guest speaker presentations organised on a monthly basis by the South Australian Committee. The meetings are held on the third Friday of each month from 4:30pm to 6:00pm. For further information, contact ATUG State Secretary, Peter Hamilton, on (08) 344 6743.

The Value of ATUG Membership

Can you afford to be without the up-to-date knowledge, provided by ATUG membership, on the evolution of convergence in the media, communications and computer industries?

As a member of ATUG you receive:

- ⇒ One year's free subscription to the monthly ATUG newsletter, 'Newsbrief';
- ⇒ One year's free subscription to *Australian Communications*;
- ⇒ Discount rates to the telecommunications event of the year, ATUG's annual four-day conference and exhibition, ATUG '95;
- ⇒ Discount rates to ATUG seminars and workshops;
- ⇒ Free admittance to ATUG special member meetings;
- ⇒ Free assistance on matters concerning telecommunications;
- ⇒ Representations of your views to Austel and the Government;
- ⇒ A forum to exchange ideas with your peers in the industry;
- ⇒ A telecommunications library service;
- ⇒ Member rates for ATUG products and services;
- ⇒ Expert legal advice from ATUG's Legal Line service; and
- ⇒ An overall saving on your next communications bill due to ATUG's lobbying activities.

If you have any queries about your ATUG membership, call ATUG's national business manager, Owen Richards, on (03) 416 1848. Members in Queensland may contact ATUG's business manager, Jane Prentice, on (07) 870 8831. Members in Western Australia may contact ATUG's WA business manager, Geoff Groves, on (09) 450 3485.

ATUG Network Management Survey

Copies of the report of ATUG's Network Management Survey are still available — but not for too long! The complete report, including quantitative data, runs to over 40 pages together with well over 60 pages of supporting information and suppliers' responses. The report is available at $125 (including postage and handling) to ATUG members; and at $200 to non-members. For copies of either the complete report or the Executive Summary, please complete the form below:

Please forward _____ (qty) copies of the ATUG Network Management Survey at $125 for each copy ($200 non-members) for which I enclose payment of: $_____

Please tick payment method (DO NOT send cash):

☐ **Cheque** ☐ **Mastercard** ☐ **Bankcard** ☐ **Visa**

Cheques should be made payable to ATUG.

Credit card number:

☐☐☐☐ ☐☐☐☐ ☐☐☐☐ ☐☐☐☐

Expires: _____ Cardholder's Signature: _____

Cardholder's Name: _____

Company: _____

Address: _____

_____

_____

OR:

☐ **Please provide a copy of the six page Executive Summary, free of charge to ATUG members only.**

Name: . . . . . .

Tel: . . . . . .

Fax: . . . . . .

Forward to:
ATUG
Network Management Survey
PO Box 357
Milsons Point, NSW 2061

August

22-23 **Telecommunications Strategic Planning**, Savoy Park Plaza, Melbourne. This intensive two-day event will provide delegates with the tools and techniques to develop strategic plans for their telecommunications networks. Topics covered include the ability to identify cost elements and build cost models for networks; analysing options for chargeback; and setting benchmarks to measure network performance. Fee: $1,295. Enquiries — AIC Training Tel: (02) 210 5777 Fax: (02) 221 7773

16-17 **GSM — Profiting from Digital Mobile Communications**, The Hyatt Kingsgate, Sydney. This conference will update delegates on the latest state-of-play and technical developments in GSM implementation. Speakers will look at the Government's planned phase-out of the AMPS network, the latest information on EMI/EMC interference, how to market and distribute GSM handsets and boost subscriber numbers, and ways to take advantage of new opportunities in value added services. Fee: $1,395. Enquiries — IIR Conferences Tel: (02) 954 5844 Fax: (02) 959 4684.

22-23 **Teleconferencing,** Hotel Inter-Continental, Sydney. With the worldwide visual services market predicted be worth $US3.2 billion in 1997, communications managers need to keep abreast of the developments in this fast-changing field. This conference features case studies from leading experts from the corporate, government and education sectors, and there will also be a workshop on implementing and managing videoconferencing, Fee: two-day conference $1,395; workshop only $495; both $1,795. Enquiries — AIC Conferences Tel: (02) 210 5777 Fax: (02) 221 7773

24-25 **Service Delivery Over the Last Mile of the Information Superhighway**, Ritz Carlton Double Bay, Sydney. This event focuses on the local loop in the information superhighway. Speakers will look at competing technologies, discuss government initiatives in the US and Europe, and give and international perspective on policy developments. There will also be a special Interim Report from Australia's Broadband Services Group and an update on the work of the Copyright Convergence Group. Fee: $1,395. Enquiries — IBC Conferences Tel: (02) 319 3755 Fax: (02) 699 3901.

September

5-6 **Cable TV and Interactive Services**, Hotel Nikko, Potts Point. This conference aims to provide the most up-to-date and comprehensive information on all aspects of cable television and interactive services. It features a panel of experts from both Australia and overseas, and will address such issues as: national policy directions; the role of cable TV in the MDS and satellite gateways; the plans of Australia's major telecommunications players; ADSL; video-on-demand; cable telephony; investment opportunities and strategic alliances. The program will also deal with copyright and other legal issues, and look at recent advances in marketing and technology. Fee: $1,395. Enquiries — IIR Conferences Tel: (02) 954 5844 Fax: (02) 959 4684.

13-14 **Global Telecoms '94**, Le Meridien, London. This international event allows delegates to choose from over 40 different sessions, tailoring their programme to their needs. Day One will comprise an introductory briefing on Worldwide Regulation Strategy, covering regulation issues at national, European and international levels. Days Two and Three will feature speakers on topics such as investment opportunities, outsourcing, competition in the local loop, market characteristics in Asia, and new opportunities in Latin America and Eastern Europe. Fee: three-day forum £1,404.13; two-day conference £1,051.63; one-day Regulation briefing £581.63. Enquiries — IIR Ltd Tel: +44 71 412 0141 Fax: +44 71 412 0145.

25-28 **LETA 94**, Adelaide Convention Centre, Adelaide. This conference, Learning Environment Technology Australia, will explore, explain and demonstrate the effects of new and forthcoming technologies on learning and the way it is undertaken. As well as teaching issues there will be a focus on building design and adaptation to meet new requirements. The event features a wide range of speakers from both Australia and overseas. Enquiries — Tel: (08) 226 1266 Fax: (08) 226 1583.

26-27 **Communications Research Forum 1994**, Hyatt Kingsgate, Sydney. This event, targeted at the research community, looks at current and recent developments relevant to emerging public policy issues in the communications arena. The focus of the two-day conference will be on telecommunications, broadcasting and the radio spectrum, looking specifically at policy-relevant work into economic, sociological, cultural and legal implications. Enquiries — Vickie Richardson Tel: (06) 274 6846 Fax: (06) 274 6816.

October

25-27 **Australian International Virtual Reality Conference**, World Congress Centre, Melbourne. Australia's first virtual reality conference will be hosted by software supplier Integra, and will feature key speakers and workshops to discuss issues in-depth and offer hands-on sessions. An exhibition will run concurrently with the conference, and will feature both commercial and entertainment-based VR applications. Enquiries — AUSI-VR'94 Conference and Exhibition Tel: (08) 234 1145 Fax: (08) 234 1148.

25-28 **Communications India**, Pragati Maidan, New Delhi. This trade show will display a full range of telecommunications hardware and software in the areas of voice, data, visual display and radio. An associated technical conference will cover cellular mobile, networking and rural telephony. Enquiries — Exhibitions India Tel: +91 11 462 2710 Fax: +91 11 463 3506.

25-27 **Microwaves '94**, Wembly Conference & Exhibition Centre, London. This conference brings together Europe's largest gathering of microwave experts engaged in R&D for telecommunications and civil and military systems. The program covers the entire microwave and millimetre wave spectrum, with supporting RF and software technologies. Enquiries — Nexus Business Communications Tel: +44 322 660 070 Fax: +44 322 667 633.

November

5-10 **MANTECH '94**, Pragati Maiden, New Delhi. This international seminar and exhibition on manufacturing technologies is organised by the Federation of Indian Chambers of Commerce and Industry. It provides participating countries with the opportunity to present state-of-the-art technologies in manufacturing, communications and electronics to Indian business, and coincides with the annual conference of the International Chamber of Commerce, which will attract delegates from 130 member countries. Enquiries — Adesh Goel, PN International Tel: (02) 328 1948 Fax: (02) 327 1975.

December

5-7 **ATNAC '94**, Hilton Hotel, Melbourne. This event incorporates the Multimedia Communications, Applications and Technology Workshop, the Australia Broadband Switching & Services Symposium, and the Australian Teletraffic Research Seminar. The theme of the conference will be 'Meeting User Needs,' and keynote speakers include Phillip Dodds, President of the Interactive Multimedia Association of the US and Leonardo Chiariglione, convenor of MPEG. The registration fee for the three-day conference is $450. Enquiries — Margaret Keegel Tel: (03) 903 2808 Fax: (03) 903 2805.

January 1995

22-26 **PTC'95**, Sheraton Waikiki Hotel, Honolulu. The 17th Annual Pacific Telecommunications Conference has the theme of 'Convergence — Closing the Gap,' and will explore the convergence of technologies and the change this is bringing about in organisations, countries and industry sectors. Enquiries — PTC Tel: +1 808 941 3789 Fax: +1 808 944 4874.

February

20-23 **InterComm 95**, Exhibition Centre, Vancouver. This year's conference has the umbrella theme 'Global Communications in the 21st Century.' Special topics under discussion include advances in wireless technology, global and corporate information networks, and cabling issues for voice, data and image. A large trade exhibition will accompany the conference. Enquiries — Tel: +1 604 669 1090 Fax: +1 604 682 5703.

March

8-15 **CeBIT '95**, Hannover Fairground, Germany. Germany's famous technical trade show will feature close on 6,000 exhibitors in 1995, in many fields, including: information technology, telecommunications, software development, computer-integrated manufacturing, office equipment, banking technologies and research services. Enquiries — German Australian Chamber of Industry and Commerce Tel: (02) 261 3982 Fax: (02) 267 3807.

20-22 **The 1995 Pan Asia Satellite and Cable Television Conference and Exhibition**, Hong Kong Convention and Exhibition Centre, Hong Kong. Now in its fifth year, this event is focused on the business opportunities in cable service provision throughout the Asia region, the world's fastest growing market for cable television and related services. There will also be an exhibtion accompanying the convention featuring a range of the latest hardware and software. Fee: $US1,795 (if registered before September 1). Enquiries — AIC Conferences Tel: +852 520 1481 Fax: +852 866 7340.

April

15-18 **Midcom '95**, Abu Dhabi International Exhibition Centre, United Arab Emirates. Midcom '95 will be the largest and most comprehensive exhibition of communications equipment and services ever staged in the Middle East region. A conference on communications issues will accompany the exhibition. Enquiries — James McGee, Infocentre International, Dubai Tel: +971 4 310 551 or US agent Chuck Heath Tel: +1 215 751 9750.

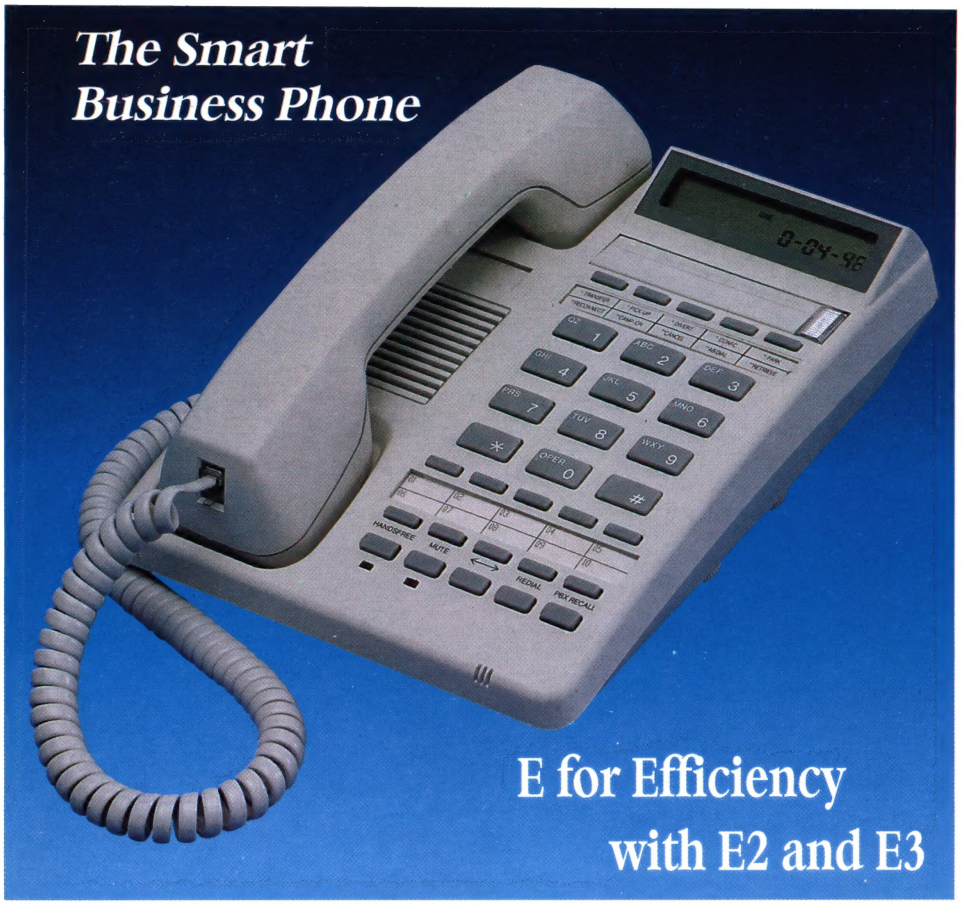

IQTEL E3

- Liquid crystal display
- Clock
- Handsfree / conference operation
- Adjustable earpiece and handsfree speaker volume
- Line powered
- Memory capacity 20 Telephone numbers or PABX commands including 10 numbers with non-erasable Eeprom insertion
- Soft touch number key pad
- Colour Grey

IQTEL E2

On Hook / Monitor Dialling Version

The following feature is optional

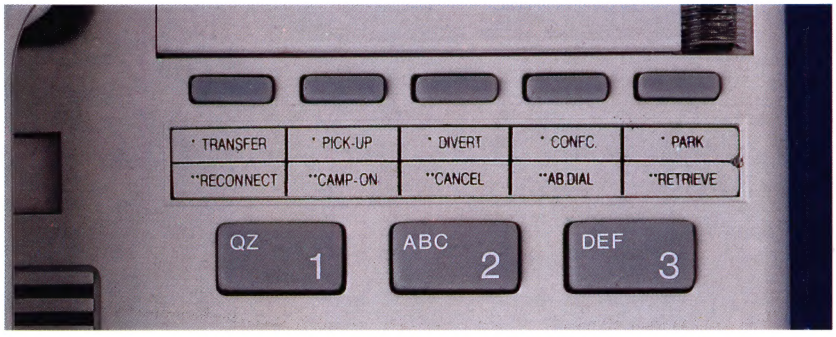

These phones not only look good, they will help you through the business day.

Easy to identify, easy to use PABX or carrier network facilities.

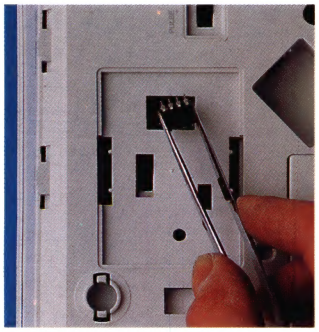

We fit this optional, specially masked Eeprom, permanently pre- programmed with your PABX or carrier network facilities

Maybe it is just checking the time or a telephone number being dialled from the display, talking handsfree or dialling on hook, phoning a colleague with one key touch or instructing the PABX or carrier network facilities to "call forward" your calls by one touch of a memory key. Whatever your job is, you will find E2 or E3 a very dependable workmate.

E2 and E3 can be purchased with customised pre-programming of PABX or network services such as Centel/Customnet/Centrex facilities as illustrated on the left of this page. The pre-programming is on a permanent basis, and will result in all staff being able to easily identify and use the PABX or Carrier facilities. The training time of staff in telephone feature use can be greatly reduced or even eliminated with E2 and E3.

E2 and E3 are not only business phones, they are ideal for home use when purchased with Telecom's "Easy Call" facilities pre-programmed into Eeprom memory for one touch activation.

IQ INTERQUARTZ® – The Manufacturers and Distributors

INTERQUARTZ (A'ASIA) PTY. LTD **Melbourne** Phone: (03) 457 5000 Fax: (03) 458 4049 **Sydney** Phone: (02) 875 4011 Fax: (02) 875 4146 **Brisbane** Phone: (07) 808 8008 Fax: (07) 808 8080

A PABX so advanced it actually talks to you.

These days, most PABX systems have a host of helpful features. Unfortunately, most people can never remember how to use all of them. That's why the new Alcatel 4300 has "Voice Prompt". It actually tells you how to use the system. For example, if you dial a busy extension, Voice Prompt will tell you a number to press to ensure a call back, or how to be automatically connected when the line is free. Of course, this is only one of many useful Voice Prompts. And Voice Prompt is only one of many features. Alcatel Australia. We're Australia's largest communications company, because we know how to compete.

ALCATEL
AUSTRALIA

For more information, send to Alcatel Australia, P.O. Box 488 North Sydney 2060 or fax to 1 800 670 015

Name Title Company

Address Postcode

☐ Send me product literature ☐ Call me to arrange a demonstration.

LBCM ALC 0119